



RESEARCH ARTICLE

Geometric approaches to matrix normalization and graph balancing

Tom Needham¹ and Clayton Shonkwiler²

¹Department of Mathematics, Florida State University, 32306 Tallahassee, FL, USA; E-mail: tneedham@fsu.edu.

²Department of Mathematics, Colorado State University, 80523 Fort Collins, CO, USA;

E-mail: clayton.shonkwiler@colostate.edu (corresponding author).

Received: 24 May 2024; Revised: 29 August 2024; Accepted: 29 August 2024

2020 Mathematical Subject Classification: Primary – 15B57; Secondary – 65K10, 90C26, 05C50

Abstract

Normal matrices, or matrices which commute with their adjoints, are of fundamental importance in pure and applied mathematics. In this paper, we study a natural functional on the space of square complex matrices whose global minimizers are normal matrices. We show that this functional, which we refer to as the non-normal energy, has incredibly well-behaved gradient descent dynamics: despite it being nonconvex, we show that the only critical points of the non-normal energy are the normal matrices, and that its gradient descent trajectories fix matrix spectra and preserve the subset of real matrices. We also show that, even when restricted to the subset of unit Frobenius norm matrices, the gradient flow of the non-normal energy retains many of these useful properties. This is applied to prove that low-dimensional homotopy groups of spaces of unit norm normal matrices vanish; for example, we show that the space of $d \times d$ complex unit norm normal matrices is simply connected for all $d \geq 2$. Finally, we consider the related problem of balancing a weighted directed graph – that is, readjusting its edge weights so that the weighted in-degree and out-degree are the same at each node. We adapt the non-normal energy to define another natural functional whose global minima are balanced graphs and show that gradient descent of this functional always converges to a balanced graph, while preserving graph spectra and realness of the weights. Our results were inspired by concepts from symplectic geometry and Geometric Invariant Theory, but we mostly avoid invoking this machinery and our proofs are generally self-contained.

1. Introduction

A matrix A is called *normal* if it commutes with its conjugate transpose: $AA^* = A^*A$. The set of $d \times d$ complex normal matrices, which we denote as $\mathcal{N}_d \subset \mathbb{C}^{d \times d}$, is a fundamental object in linear algebra; for example, the Spectral Theorem characterizes \mathcal{N}_d as the set of unitarily diagonalizable matrices:

$$\mathcal{N}_d = \{UDU^* \mid D \in \mathcal{D}_d, U \in U(d)\},$$

where $\mathcal{D}_d \subset \mathbb{C}^{d \times d}$ is the set of diagonal matrices and $U(d)$ is the group of unitary matrices. Moreover, normal matrices are especially well-behaved from a numerical analysis perspective. Indeed, the Bauer–Fike Theorem [5] implies that the eigenvalues of a normal matrix are Lipschitz stable under perturbations, which motivates the approximation of transfer matrices by normal matrices in classical control theory [15, 16]. In the literature on dynamics on complex networks, it has also been observed that directed networks whose weighted adjacency matrices are not normal exhibit distinctive dynamical features which can confound classical spectral methods [2, 3, 45]. Based on these considerations, the *closest*

normal matrix problem – that is, the problem of finding a closest point in \mathcal{N}_d to an arbitrary matrix in $\mathbb{C}^{d \times d}$ – has been thoroughly studied [22, 25, 50, 52].

This paper studies \mathcal{N}_d from a geometric perspective, with a view toward optimization tasks such as the closest normal matrix problem. Our results are largely derived from the simple observation that \mathcal{N}_d is the set of global minima of the function

$$E : \mathbb{C}^{d \times d} \rightarrow \mathbb{R}, \quad \text{defined by} \quad E(A) = \|AA^* - A^*A\|^2, \quad (1.1)$$

where $\|\cdot\|$ is the Frobenius norm; that is,

$$\|B\|^2 = \sum_{i,j=1}^d b_{ij}^2, \quad \text{for} \quad B = (b_{ij})_{i,j=1}^d \in \mathbb{C}^{d \times d}.$$

Despite the fact that the function E , which we refer to as the *non-normal energy*, is not quasiconvex (see Remark 2.5), it is surprisingly well-behaved from an optimization perspective: we prove in Theorem 2.3 that the only critical points of E are normal matrices, hence gradient descent along E gives an approximate solution to the closest normal matrix problem. We derive several related results, which are described in more detail below in Section 1.1; in short, we show that gradient descent preserves interesting features of the initializing matrix, such as its spectrum or the realness of its entries. We also consider the restriction of non-normal energy to the space of matrices with unit Frobenius norm and show that its gradient flow is also quite well-behaved. This has immediate topological implications, as we explain in more detail in Section 1.1.

The properties of the non-normal energy which we exploit in this paper are predictable from a high-level perspective: E is the squared norm of a momentum map associated to a Hamiltonian action of $SU(d)$ on $\mathbb{C}^{d \times d}$ (see Proposition 2.2). This terminology comes from the field of symplectic geometry, where the behavior of functions of this form is well-understood [37, 40]. Our work in this paper is heavily inspired by Mumford's *Geometric Invariant Theory (GIT)* [44] (see [55] for a nice introduction or [43] for applications to similar matrix optimization problems) and Kirwan's work relating GIT and symplectic geometry [37]; there are also strong connections to Ness's paper [49]. One of our goals in writing this paper was to make our arguments – especially the fundamentally elementary ones – as accessible as possible, so we have mostly avoided explicitly invoking GIT in what follows, but it was very much on our minds as we were working on this paper. Connections to GIT and symplectic ideas are explained throughout.

As our results on \mathcal{N}_d are rooted in powerful general theory, it should not be surprising that our techniques are more broadly applicable. Indeed, we also apply our geometric approach to the *graph balancing problem*: given a weighted, directed graph \mathcal{G} , one wishes to determine a new set of edge weights which balances the graph in the sense that the weighted in-degree is the same as the weighted out-degree at each node. If the latter condition is met, we say that the graph is *balanced*. An example of our gradient flow-based approach to graph balancing, as is described below, is shown in Figure 1. This problem is natural from an applications perspective; for example, in the case that the underlying graph represents a road network and that the weights are roadway capacities, that the graph is balanced corresponds to the feasibility of traffic flow through all intersections. As such, the graph balancing problem is well-studied in the operations research literature [26, 33, 51].

Representing a graph \mathcal{G} on d nodes by a matrix $A \in \mathbb{C}^{d \times d}$ containing the square roots of the entries of the weighted adjacency matrix of \mathcal{G} , the balanced graphs are exactly the global minima of the *unbalanced energy* function,

$$B : \mathbb{C}^{d \times d} \rightarrow \mathbb{R}, \quad \text{defined by} \quad A \mapsto \|\text{diag}(AA^* - A^*A)\|^2, \quad (1.2)$$

where diag is the linear map which zeros out all off-diagonal entries. The unbalanced energy is similar in structure to the non-normal energy – in fact, it is also the squared norm of a momentum map – and




Figure 1. Balancing a graph, starting at top left with a random weighted, directed multigraph with 6 vertices and 15 edges and ending with a balanced graph with the same edges and vertices on the bottom right. The thickness of each edge is proportional to its weight and the time parameter is logarithmic in the number of iterations of gradient descent. Two features of interest: different edges have activity in different timeframes (compare the two edges connecting the bottom-right vertex to the top-center vertex), and the weight of an edge can be nonmonotone as a function of time (e.g., the left-most edge or the edge connecting the top-right vertex to the central vertex).

we derive similar results regarding its gradient flow. We show in Theorem 4.1 that the critical points of B are exactly the balanced matrices and refine this result to show that gradient flow preserves geometric features of the underlying graph. We describe these results more precisely in the following subsection.

1.1. Main contributions and outline

We now summarize our main results in more detail.

- **Gradient flow of non-normal energy:** Section 2 considers properties of the non-normal energy (1.1), with a focus on properties of its gradient descent dynamics in relation to normal matrices. Although the non-normal energy is not convex (Remark 2.5), we show in Theorem 2.3 that the only critical points of E are normal matrices; that is, its global minima. It follows easily that its gradient descent has a well-defined limiting normal matrix for every choice of initial conditions; we additionally show in Theorem 2.6 that the gradient descent trajectories of the non-normal entries preserve spectra and realness of matrix entries. We derive estimates of the distance traveled under gradient flow, which give new interpretations of concepts in the literature on the closest normal matrix problem (Corollary 2.9 and Proposition 2.11).
- **Restriction to unit norm matrices and topological consequences:** In Section 3, we consider the restriction of the non-normal energy to the space of matrices with unit Frobenius norm. We prove in Theorem 3.1 that if gradient descent is initialized at a non-nilpotent unit norm matrix, then it converges to a normal matrix, and that if the initialization has real entries then so does its limit. As an application, we show that the low-dimension homotopy groups of the spaces of complex and real unit norm normal matrices vanish in Theorem 3.9 and Theorem 3.13, respectively. In particular, the space of $d \times d$ unit norm complex normal matrices is connected for all d and simply connected for $d \geq 2$, whereas the space of unit norm real normal matrices is connected for $d \geq 2$ and simply connected for $d \geq 3$.

- **Graph balancing via unbalanced energy:** The unbalanced energy (1.2) and its applications to graph balancing are studied in Section 4. Theorem 4.1 shows that the only critical points of the unbalanced energy are its global minima; that is, matrices representing balanced digraphs. Gradient descent converges to a balanced digraph representation, and we show in Theorem 4.3 that it preserves spectra and realness of entries. Moreover, this theorem shows that if the entries of a real matrix are positive then this property is also preserved, and that if an entry in the initial matrix is zero then it stays zero along the gradient descent path – in terms of graphs, gradient descent does not create any edges that were not present at initialization. We also consider the restriction of the unbalanced energy to unit norm matrices (which represent digraphs with a fixed total edge capacity) and derive similar useful properties of its gradient flow in Theorem 4.5. Finally, we observe in Theorem 4.10 that the spaces of complex and real balanced unit norm matrices are homotopy equivalent to spaces of real and complex normal matrices, respectively.

2. Normal matrices and optimization

Recall from the introduction that the *non-normal energy* $E : \mathbb{C}^{d \times d} \rightarrow \mathbb{R}$ is the function

$$E(A) = \|AA^* - A^*A\|^2 = \|[A, A^*]\|^2.$$

Throughout this paper we use $[\cdot, \cdot]$ to denote the matrix commutator: $[A, B] = AB - BA$.

The goal of this section is to derive properties of the gradient descent dynamics of E . In particular, we will show that we can normalize any square matrix by sending it to its limit under the negative gradient flow of E .

2.1. Background

The map E has a long history in the problem of finding the closest normal matrix to a given matrix, going back at least to Henrici [30], who proved the following:

Proposition 2.1 (Henrici [30]). *For any $A \in \mathbb{C}^{d \times d}$,*

$$\inf_{M \in \mathcal{N}_d} \|A - M\| \leq \left(\frac{d^3 - d}{12} E(A) \right)^{1/4}.$$

In other words, the distance from A to \mathcal{N}_d is bounded above by a quantity proportional to $E(A)^{1/4}$. One virtue of this estimate is that $E(A)$ is relatively easy to compute.

We now give an interpretation of E in terms of symplectic geometry, where we consider $\mathbb{C}^{d \times d} \approx \mathbb{C}^{d^2}$ as a symplectic manifold with its standard symplectic structure. This interpretation is not necessary for most of the paper, and is mainly included for context. As such, we give a somewhat informal treatment and avoid explicit definitions of any of the standard terminology from symplectic geometry. In our previous papers, we give short and elementary overviews of the necessary concepts from symplectic geometry, with a view toward understanding similar spaces of structured matrices (e.g., spaces consisting of unit norm tight frames); we refer the reader to [46, Section 2] and [47, Section 2.1] for more in-depth exposition.

Consider the action of the unitary group $SU(d)$ on $\mathbb{C}^{d \times d}$ by conjugation. Let $\mathfrak{su}(d)$ denote the Lie algebra of $SU(d)$ – that is, the traceless, skew-Hermitian $d \times d$ matrices – and let $\mathfrak{su}(d)^*$ denote its dual. It will be convenient to identify $\mathfrak{su}(d)^*$ with the space $\mathcal{H}_0(d)$ of $d \times d$ traceless Hermitian matrices via the isomorphism

$$\begin{aligned} \mathcal{H}_0(d) &\rightarrow \mathfrak{su}(d)^* \\ Y &\mapsto (X \mapsto \frac{i}{2} \operatorname{Tr}(XY)). \end{aligned}$$

Then we have the following interpretation of E .

Proposition 2.2. *The conjugation action of $\mathrm{SU}(d)$ on $\mathbb{C}^{d \times d}$ is Hamiltonian, with momentum map*

$$\mu : \mathbb{C}^{d \times d} \rightarrow \mathcal{H}_0(d) \approx \mathfrak{su}(d)^* \quad (2.1)$$

$$A \mapsto [A, A^*]. \quad (2.2)$$

The non-normal energy E is therefore the squared norm of a momentum map.

We omit the proof of Proposition 2.2, which is a straightforward calculation. In light of this result, one should expect the non-normal energy to have nice properties – see, for example, work of Kirwan [37] and Lerman [40]. However, the specific properties of E (and related functions) that we derive below do not follow directly from the general theory.

2.2. Critical points of non-normal energy

Obviously, the global minima of the non-normal energy E are exactly the normal matrices. In fact, we now show that these are the only critical points. Throughout the paper, we use $\langle \cdot, \cdot \rangle$ to denote the real part of the Frobenius inner product on $\mathbb{C}^{d \times d}$,

$$\langle A, B \rangle := \operatorname{Re} \operatorname{tr}(B^* A),$$

and we use $DF(A)$ to denote the derivative of a map $F : \mathbb{C}^{d \times d} \rightarrow \mathbb{R}$ at $A \in \mathbb{C}^{d \times d}$.

Theorem 2.3. *The only critical points of E are the global minima; that is, the normal matrices.*

Proof. We claim that

$$\nabla E(A) = -4[A, [A, A^*]]. \quad (2.3)$$

Indeed, since E is the square of a momentum map (Proposition 2.2), this follows by general principles of symplectic geometry – see, for example, [37, Lemma 6.6] or [49, Lemma 6.1]. Let us additionally give an elementary derivation of this fact. Writing $E = N \circ \mu$, where μ is the momentum map (2.1) and $N : \mathbb{C}^{d \times d} \rightarrow \mathbb{R}$ is the norm-squared map $N(A) = \|A\|^2$, we have, for any $A, B \in \mathbb{C}^{d \times d}$,

$$\begin{aligned} \langle \nabla E(A), B \rangle &= DE(A)(B) = DN(\mu(A)) \circ D\mu(A)(B) \\ &= \langle \nabla N(\mu(A)), D\mu(A)(B) \rangle = \langle D\mu(A)^\vee \nabla N(\mu(A)), B \rangle, \end{aligned}$$

where we use $D\mu(A)^\vee$ to denote the adjoint of $D\mu(A)$ with respect to the inner product $\langle \cdot, \cdot \rangle$. It follows that

$$\nabla E(A) = D\mu(A)^\vee \nabla N(\mu(A)).$$

A straightforward calculation then shows that the adjoint is given by the formula

$$D\mu(A)^\vee(C) = [C + C^*, A]. \quad (2.4)$$

It is also easy to show that $\nabla N(C) = 2C$, so we conclude that

$$\nabla E(A) = [2\mu(A) + 2\mu(A)^*, A] = -4[A, [A, A^*]].$$

Therefore, we have a critical point of E exactly when

$$0 = [A, [A, A^*]];$$




Figure 2. The graph of E restricted to the collection of real matrices of the form $\begin{bmatrix} 0 & x \\ y & 0 \end{bmatrix}$.

that is, when A and $[A, A^*]$ commute. By Jacobson's Lemma (stated below as Lemma 2.4), this implies that $[A, A^*]$ is nilpotent. But $[A, A^*]$ is Hermitian, so it is nilpotent if and only if it is the zero matrix, which happens precisely when A is normal.¹ \square

Lemma 2.4 (Jacobson [34]; see also [36]). *If A and B are $d \times d$ matrices over a field of characteristic 0 and A commutes with $[A, B]$, then $[A, B]$ is nilpotent.*

Remark 2.5. Theorem 2.3 might lead one to suspect that E is convex, but it is not. To see this, consider the normal matrices

$$A_0 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \quad \text{and} \quad A_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Since they are normal, $E(A_0) = 0 = E(A_1)$. However,

$$E((1-t)A_0 + tA_1) = 32t^2(1-t)^2 > 0$$

for all $0 < t < 1$, so the interior of the line segment connecting A_0 and A_1 consists entirely of non-normal matrices, and hence E is not even quasiconvex, let alone convex. See Figure 2. Of course, we can pad A_0 and A_1 by zeros to get an analogous example for any $d > 2$.

On the other hand, Theorem 2.3 shows that E is an *invex function*. Recall that, as first defined by Hanson [27] (later named by Craven [13]), a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is *invex* if there exists a function $\eta : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that

$$f(x) - f(u) \geq \langle \eta(x, u), \nabla f(u) \rangle \quad x, u \in \mathbb{R}^n.$$

A theorem of Craven and Glover [14] (see also [6]) says that a function is invex if and only if its critical points are all global minima; hence, E is invex.

2.3. Gradient flow of non-normal energy

Consider the negative gradient flow $\mathcal{F} : \mathbb{C}^{d \times d} \times [0, \infty) \rightarrow \mathbb{C}^{d \times d}$ defined by

$$\mathcal{F}(A_0, 0) = A_0, \quad \frac{d}{dt} \mathcal{F}(A_0, t) = -\nabla E(\mathcal{F}(A_0, t)). \quad (2.5)$$

We pause here to note that there is a substantial history of applying flows like (2.5) to problems in numerical linear algebra, going back at least to Rutishauser's work on the LU decomposition [53, Section 11]. See Chu's survey [12] for an introduction to this circle of ideas.

¹The equivalence of A being normal and A commuting with $[A, A^*]$ appears as #73 in Elsner and Ikramov's list [19]; they attribute it to [42, 4.28.5].

Since E is a real polynomial function on the real vector space $\mathbb{C}^{d \times d}$, the gradient flow cannot have limit cycles or other bad behavior [41], so Theorem 2.3 implies that, for any $A_0 \in \mathbb{C}^{d \times d}$, the limit $A_\infty := \lim_{t \rightarrow \infty} \mathcal{F}(A_0, t)$ of the gradient flow is well-defined and normal.

From (2.3), we see that

$$\begin{aligned} \nabla E(A) &= -4[A, [A, A^*]] = -4(A[A, A^*] - [A, A^*]A) \\ &= 4 \left(\frac{d}{d\epsilon} \bigg|_{\epsilon=0} e^{\epsilon[A, A^*]} A e^{-\epsilon[A, A^*]} \right) = 4 \left(\frac{d}{d\epsilon} \bigg|_{\epsilon=0} e^{\epsilon[A, A^*]} \cdot A \right). \end{aligned} \quad (2.6)$$

Since $[A, A^*]$ is traceless, $e^{\epsilon[A, A^*]} \in \mathrm{SL}_d(\mathbb{C})$ for any ϵ , so the negative gradient flow lines $\mathcal{F}(A_0, t)$ produced by any A_0 stay within the conjugation orbit of A_0 . In particular, A_∞ must have the same eigenvalues as A_0 . Since real matrices are invariant under gradient flow, we have thus proved:

Theorem 2.6. *For any $A_0 \in \mathbb{C}^{d \times d}$, the matrix $A_\infty = \lim_{t \rightarrow \infty} \mathcal{F}(A_0, t)$ exists, is normal, and has the same eigenvalues as A_0 . Moreover, if A_0 is real, then so is A_∞ .*

Remark 2.7. This theorem implies that the gradient flow (2.5) is an *isospectral flow* [17, 56, 57, 58]. While it does not quite fit into the classical framework of isospectral flows except in the trivial case that A_0 is Hermitian, we caution readers interested in implementing (2.5) that standard numerical ODE methods fail in the classical case [10], and there is no reason to think that our flow won't pose similar numerical issues.

If A_0 and A_∞ are as in Theorem 2.6 and $\lambda_1, \dots, \lambda_d$ are their common eigenvalues, then the normality of A_∞ implies that

$$\|A_\infty\|^2 = \sum_{i=1}^d |\lambda_i|^2.$$

This immediately implies the following corollary.

Corollary 2.8. *If A_0 is non-nilpotent, then its gradient flow (2.5) is bounded away from zero. On the other hand, if A_0 is nilpotent, then the limit of gradient flow A_∞ is the zero matrix.*

A widely used statistic for describing the extent to which a matrix is non-normal is the *Henrici departure from normality* [30]. For a matrix $A \in \mathbb{C}^{d \times d}$ with eigenvalues λ_i , this is the quantity²

$$\mathrm{Hen}(A) = \|A\|^2 - \sum_{i=1}^d |\lambda_i|^2.$$

Corollary 2.9. *Let $A_0 \in \mathbb{C}^{d \times d}$ and let A_∞ be its limit under the gradient flow (2.5). The change in scale along gradient flow is equal to Henrici departure from normality,*

$$\|A_0\|^2 - \|A_\infty\|^2 = \mathrm{Hen}(A_0).$$

2.4. Bound on the distance to the limit of gradient flow

We now show that A_∞ is not too much further from A_0 than the closest normal matrix, despite the fact that A_∞ preserves features (spectrum, realness) that the closest normal matrix may not. We do so by a standard argument starting from a Łojasiewicz inequality.

²Note that some authors, including Henrici, define the departure from normality to be the square root of this quantity.

Since E is the squared norm of a momentum map (Proposition 2.2), a result of Fisher [20] gives us the desired inequality³:

Proposition 2.10 (Fisher [20, Theorem 4.7]). *There exist constants $\epsilon, c > 0$ so that for all $A \in \mathbb{C}^{d \times d}$ with $E(A) < \epsilon$,*

$$\|\nabla E(A)\| \geq c E(A)^{3/4}.$$

Now we follow a standard argument (see, e.g., Lerman [40]) to get bounds on the distance from A_0 to A_∞ . Certainly this distance is no larger than the length of the gradient flow path:

$$\|A_0 - A_\infty\| \leq \int_0^\infty \left\| \frac{d}{dt} \mathcal{F}(A_0, t) \right\| dt = \int_0^\infty \|\nabla E(\mathcal{F}(A_0, t))\| dt. \quad (2.7)$$

So long as $E(\mathcal{F}(A_0, t)) < \epsilon$,

$$\begin{aligned} -\frac{d}{dt} (E(\mathcal{F}(A_0, t)))^{1/4} &= -\frac{1}{4} E(\mathcal{F}(A_0, t))^{-3/4} D E(\mathcal{F}(A_0, t)) (-\nabla E(\mathcal{F}(A_0, t))) \\ &= \frac{1}{4} E(\mathcal{F}(A_0, t))^{-3/4} \|\nabla E(\mathcal{F}(A_0, t))\|^2 \geq \frac{c}{4} \|\nabla E(\mathcal{F}(A_0, t))\|, \end{aligned}$$

where the last inequality follows since Proposition 2.10 implies $E(\mathcal{F}(A_0, t))^{-3/4} \|\nabla E(\mathcal{F}(A_0, t))\| \geq c$.

Combining this with (2.7) yields:

$$\|A_0 - A_\infty\| \leq \int_0^\infty \|\nabla E(\mathcal{F}(A_0, t))\| dt \leq -\frac{4}{c} \int_0^\infty \frac{d}{dt} (E(\mathcal{F}(A_0, t)))^{1/4} dt = \frac{4}{c} E(A_0)^{1/4}.$$

Therefore, we have proved:

Proposition 2.11. *There exist constants $\epsilon, c > 0$ so that, if $E(A_0) < \epsilon$, then*

$$\|A_0 - A_\infty\| \leq \frac{4}{c} E(A_0)^{1/4}.$$

Comparing to the Henrici estimate (Proposition 2.1), we see that, at least when $E(A_0)$ is small, the normal matrix A_∞ we get by doing gradient descent is not much further from A_0 than the closest normal matrix is, even though A_∞ has the same spectrum as A_0 and is real if A_0 is.

Remark 2.12. The closest normal matrix to a given $A_0 \in \mathbb{C}^{d \times d}$ can be computed explicitly by Ruhe's algorithm [52],⁴ but the actual closest normal matrix does not have the same spectrum as A_0 and may be complex even if A_0 is real (see discussion in Chu [11] and Guglielmi and Scalone [25]). This suggests that the gradient descent approach to finding a nearby normal matrix may be useful in situations where one is interested in preserving structural properties of the initialization. These observations are borne out by numerical experiments, and indeed A_∞ gets relatively closer to the closest normal matrix when A_0 is almost normal to begin with: see Figure 3.

3. Unit norm normal matrices

We have seen in Corollary 2.9 that the gradient flow of E does not preserve the Frobenius norm. If we want a flow that preserves the norm, we should consider the restriction of E to the space \mathcal{U}_d of $d \times d$ matrices with Frobenius norm 1. Geometrically, \mathcal{U}_d is just the $(2d^2 - 1)$ -dimensional unit sphere in $\mathbb{C}^{d \times d}$.

³Similar results appear in Neeman [48, Theorem A.1], Woodward [60, Lemma B.0.6], and Lerman [40]; both Woodward and Lerman credit Duistermaat with proving a version of this result in unpublished work, as do Mumford, Fogarty, and Kirwan [44, p. 166, footnote 58].

⁴Note that, as Higham points out [31], there is a missing minus sign before the determinant in the definition of θ in step 2 of the published version of Ruhe's Algorithm J.




Figure 3. Left: We generated 10,000 initial matrices $A_0 \in \mathbb{C}^{20 \times 20}$ by letting the real and imaginary parts of each entry be drawn from a standard Gaussian and then normalizing so that A_0 has Frobenius norm 1. We computed the closest normal matrix \hat{A} using Ruhe's algorithm [52] and $A_\infty = \lim_{t \rightarrow \infty} \mathcal{F}(A_0, t)$ using a very simple gradient descent with fixed step sizes, and then plotted the point $(\|\hat{A} - A_0\|^2, \|A_\infty - A_0\|^2)$. The ratios $\frac{\|A_\infty - A_0\|^2}{\|\hat{A} - A_0\|^2}$ were all in the interval $[1.028, 1.161]$. Center: The same computations and visualization, except the initial matrices A_0 were all 20×20 real matrices. In this case the $\frac{\|A_\infty - A_0\|^2}{\|\hat{A} - A_0\|^2}$ were all in the interval $[1.023, 1.196]$. Right: The same computations and visualization, but with nearly normal initial matrices $A_0 \in \mathbb{C}^{20 \times 20}$. More precisely, we generated $B \in \mathbb{C}^{20 \times 20}$ by normalizing a matrix of standard complex Gaussians, found the closest normal matrix \hat{B} , then added an $\mathcal{N}(0, 0.0075)$ random variate to the real and complex parts of each entry of \hat{B} , and let A_0 be the normalization of this matrix, so that A_0 has Frobenius norm 1 and is already close to being normal. In this case the $\frac{\|A_\infty - A_0\|^2}{\|\hat{A} - A_0\|^2}$ were all in the interval $[1.009, 1.036]$. In all three plots, the solid line has slope 1 and the dashed line has slope 1.3. Code for these experiments is available on GitHub [54].

Let $\bar{E} : \mathcal{U}_d \rightarrow \mathbb{R}$ be the restriction of E to \mathcal{U}_d and let $\bar{\mathcal{F}} : \mathcal{U}_d \times [0, \infty) \rightarrow \mathcal{U}_d$ be the associated gradient flow; that is,

$$\bar{\mathcal{F}}(A_0, 0) = A_0 \quad \frac{d}{dt} \bar{\mathcal{F}}(A_0, t) = -\text{grad } \bar{E}(\bar{\mathcal{F}}(A_0, t)),$$

where grad is the Riemannian gradient on \mathcal{U}_d .

3.1. Gradient flow of restricted non-normal energy

The normal matrices in \mathcal{U}_d are exactly the global minima of \bar{E} ; the goal is to show that almost every matrix in \mathcal{U}_d flows to a normal matrix under the gradient flow:

Theorem 3.1. *For any non-nilpotent $A_0 \in \mathcal{U}_d$, the matrix $A_\infty := \lim_{t \rightarrow \infty} \bar{\mathcal{F}}(A_0, t)$ exists, is normal, and has Frobenius norm 1. Moreover, if A_0 is real, then so is A_∞ .*

Remark 3.2. In GIT terms, we are looking at (a linearization of) the projective adjoint action of $\text{SL}(d)$ on $\mathbb{P}(\mathfrak{sl}(d)^*)$, and the fact that we have to assume A_0 is non-nilpotent in Theorem 3.1 is equivalent to the fact that the non-nilpotent matrices are exactly the semistable points with respect to this action [38] (see [44, Proposition 4.4]).

Since \bar{E} is a polynomial function defined on a real-analytic submanifold of Euclidean space, it will have a Łojasiewicz exponent (cf. [7, Corollary 4.2]), and hence the gradient flow will have a single limit point [41], proving the existence of A_∞ .

Since the non-nilpotent matrices form an open, dense subset of \mathcal{U}_d , Theorem 3.1 implies that almost every member of any neighborhood of a nonminimizing critical point will flow to a normal matrix; that

is, a global minimum of \bar{E} . Hence, the nonminimizing critical points of \bar{E} cannot be basins of attraction. Since \bar{E} has a Łojasiewicz exponent, an argument analogous to [1, Theorem 3] shows that all local minima must be basins of attraction. Hence we have the following corollary.

Corollary 3.3. *Every local minimum of \bar{E} must be a global minimum; that is, a normal matrix.*

We have already shown that the gradient flow of \bar{E} converges to a single limit point A_∞ . The remainder of this subsection will be devoted to proving the remaining statements of Theorem 3.1 through several supporting results. The strategy for proving the rest of the first sentence of Theorem 3.1 is to show that the gradient flow preserves non-nilpotency and that all nonminimizing critical points must be nilpotent. As with Theorem 2.6, the last sentence will follow because the real submanifold of \mathcal{U}_d is invariant under the gradient flow.

Proposition 3.4. *The intrinsic gradient of \bar{E} on \mathcal{U}_d is*

$$\text{grad } \bar{E}(A) = -4([A, [A, A^*]] + \bar{E}(A)A).$$

Proof. Geometrically, $\text{grad } \bar{E}(A)$ is the projection of the Euclidean gradient $\nabla E(A)$ onto the tangent space $T_A \mathcal{U}_d = \text{span}(\{A\})^\perp$:

$$\text{grad } \bar{E}(A) = \nabla E(A) - \langle \nabla E(A), A \rangle A.$$

We know from (2.3) that $\nabla E(A) = -4[A, [A, A^*]]$, so the fact that $[A, A^*]$ is Hermitian implies

$$\begin{aligned} \langle \nabla E(A), A \rangle &= -4\text{Re tr}([A, [A, A^*]]^* A) = -4\text{Re tr}([A, A^*]A^* A - A^*[A, A^*]A) \\ &= 4\text{Re tr}([A, A^*][A, A^*]) = 4\| [A, A^*] \|^2 = 4\bar{E}(A) \end{aligned} \quad (3.1)$$

by the linearity and cyclic invariance of trace, and the result follows. \square

Since $[A, A^*]$ is traceless, notice that

$$\text{grad } \bar{E}(A) = 4 \frac{d}{dt} \Big|_{t=0} e^{-t\bar{E}(A)} e^{t[A, A^*]} A e^{-t[A, A^*]} = 4 \frac{d}{dt} \Big|_{t=0} (e^{t[A, A^*]}, e^{-t\bar{E}(A)}) \cdot A$$

is tangent to the $\text{SL}_d(\mathbb{C}) \times \mathbb{C}^\times$ -orbit of A , where the action of $\text{SL}_d(\mathbb{C}) \times \mathbb{C}^\times$ on $\mathbb{C}^{d \times d}$ is defined by $(g, z) \cdot A := zgAg^{-1}$.

We could use this to show that the negative gradient flow preserves non-nilpotency, but extending to the limit poses challenges, so we adopt a different approach. For $A \in \mathbb{C}^{d \times d}$, define

$$s(A) := \sum_{i=1}^d |\lambda_i|^2,$$

where $\lambda_1, \dots, \lambda_d$ are the eigenvalues of A . The nilpotent matrices are precisely the vanishing locus of s .

Lemma 3.5. *For any $A \in \mathcal{U}_d$,*

$$\langle -\text{grad } \bar{E}(A), \text{grad } s(A) \rangle = 8s(A)\bar{E}(A),$$

where $\text{grad } s(A)$ is the intrinsic gradient of s in \mathcal{U}_d .

Proof. Note, first of all, that $\langle A, \text{grad } s(A) \rangle = 0$, since $\text{grad } s(A) \in T_A \mathcal{U}_d = \text{span}(\{A\})^\perp$. Therefore,

$$\begin{aligned} \langle -\text{grad } \bar{E}(A), \text{grad } s(A) \rangle &= \langle -\nabla E(A) + 4\bar{E}(A)A, \text{grad } s(A) \rangle \\ &= \langle -\nabla E(A), \text{grad } s(A) \rangle \\ &= \langle -\nabla E(A), \nabla s(A) - \langle \nabla s(A), A \rangle A \rangle \\ &= -\langle \nabla E(A), \nabla s(A) \rangle + \langle \nabla s(A), A \rangle \langle \nabla E(A), A \rangle \\ &= -\langle \nabla E(A), \nabla s(A) \rangle + 4\langle \nabla s(A), A \rangle \bar{E}(A), \end{aligned}$$

using (3.1) in the first and last equalities.

We know from (2.6) and the following sentence that $\nabla E(A)$ lies in the conjugation orbit of A . But this means that $\nabla E(A)$ must be tangent to the level set of s passing through A , since conjugation preserves eigenvalues, and hence fixes s . Therefore, $\langle \nabla E(A), \nabla s(A) \rangle = 0$ and we have shown that

$$\langle -\text{grad } \bar{E}(A), \text{grad } s(A) \rangle = 4\langle \nabla s(A), A \rangle \bar{E}(A).$$

By definition of the gradient, the inner product is a directional derivative,

$$\langle \nabla s(A), A \rangle = Ds(A)(A) = \lim_{t \rightarrow 0} \frac{s(A + tA) - s(A)}{t} = \lim_{t \rightarrow 0} \frac{(1+t)s(A) - s(A)}{t} = s(A),$$

completing the proof. \square

Proposition 3.6. *If $A_0 \in \mathcal{U}_d$ is non-nilpotent, then so is $A_t := \bar{\mathcal{F}}(A_0, t)$ for all $t \in [0, \infty)$ and so is $A_\infty := \lim_{t \rightarrow \infty} \bar{\mathcal{F}}(A_0, t)$.*

Proof. For any $A \in \mathcal{U}_d$, Lemma 3.5 implies that

$$\langle -\text{grad } \bar{E}(A), \text{grad } s(A) \rangle = 8s(A)\bar{E}(A) \geq 0.$$

Therefore, $s(A)$ must be nondecreasing along the negative gradient flow lines of \bar{E} , so $s(A_t) \geq s(A_0) > 0$ for all $t \in [0, \infty)$, and in the limit we also have $s(A_\infty) \geq s(A_0) > 0$. Hence, A_t and A_∞ must be non-nilpotent. \square

In other words, gradient flow preserves non-nilpotency, including in the limit, so we have completed the first step in our strategy for proving Theorem 3.1. We now proceed with the second step.

Proposition 3.7. *All nonminimizing critical points of \bar{E} are nilpotent.*

Proof. By Proposition 3.4, A is a critical point of \bar{E} if and only if

$$0 = [A, [A, A^*]] + \bar{E}(A)A.$$

If A is a nonminimizing critical point, then A is not normal, so $\bar{E}(A) \neq 0$ and

$$A = -\frac{1}{\bar{E}(A)}[A, [A, A^*]].$$

In other words, $A = [A, B]$ with $B = -\frac{1}{\bar{E}(A)}[A, A^*]$. But then A certainly commutes with $[A, B]$, so Jacobson's Lemma (Lemma 2.4) implies that $[A, B]$ is nilpotent. Since $A = [A, B]$, we conclude that A is nilpotent. \square

Proof of Theorem 3.1. If $A_0 \in \mathcal{U}_d$ is not nilpotent, then the limit $A_\infty = \lim_{t \rightarrow \infty} \bar{\mathcal{F}}(A_0, t)$ exists and, by Proposition 3.6, is not nilpotent. A_∞ must be a critical point of \bar{E} and, by Proposition 3.7, must be a global minimum, and hence normal. \square




Figure 4. This is the same experimental setup as in Figure 3, except that now $A_\infty = \lim_{t \rightarrow \infty} \overline{\mathcal{F}}(A_0, t)$. Left: $A_0 \in \mathbb{C}^{20 \times 20}$; all $\frac{\|A_\infty - A_0\|^2}{\|A - A_0\|^2} \in [1.060, 1.198]$. Center: $A_0 \in \mathbb{R}^{20 \times 20}$; all $\frac{\|A_\infty - A_0\|^2}{\|A - A_0\|^2} \in [1.046, 1.253]$. Right: $A_0 \in \mathbb{C}^{20 \times 20}$ is a small perturbation of a normal matrix; all $\frac{\|A_\infty - A_0\|^2}{\|A - A_0\|^2} \in [1.010, 1.031]$. In all three plots, the solid line has slope 1 and the dashed line has slope 1.3. Code for these experiments is available on GitHub [54].

It is possible to prove an analogous statement to Proposition 2.11 in this setting as well, so gradient descent of \overline{E} , even though it preserves norms and (when applicable) realness, produces a limiting normal matrix A_∞ which is not much further from A_0 than the closest normal matrix. Again, this conclusion is supported by numerical experiments: see Figure 4.

3.2. Topology of unit norm normal matrices

The space of normal matrices \mathcal{N}_d is a cone in $\mathbb{C}^{d \times d}$ and hence topologically trivial. However, the space \mathcal{UN}_d can potentially have interesting topology. Friedland [21] argues that \mathcal{UN}_d is irreducible and the quasi-variety of its smooth points is connected. However, this is not quite enough to imply that \mathcal{UN}_d is connected, since irreducible real varieties can have connected components consisting entirely of nonsmooth points (see, e.g., [9, Figure 2]). In this subsection, we show that \mathcal{UN}_d is connected and, in fact, that many of its low-dimensional homotopy groups vanish.

The key fact that we use when studying the topology of \mathcal{UN}_d is that it is closely related to the topology of the space of all non-nilpotent matrices in $\mathbb{C}^{d \times d}$. For the rest of this subsection, we use \mathcal{P}_d to denote the space of nilpotent matrices in $\mathbb{C}^{d \times d}$ and we let $\mathcal{M}_d = \mathbb{C}^{d \times d} \setminus \mathcal{P}_d$. The relationship between the topologies of \mathcal{UN}_d and \mathcal{M}_d is made precise by the following result.

Corollary 3.8. *The space \mathcal{UN}_d is a strong deformation retract of \mathcal{M}_d .*

Proof. As the function $\mu : A \mapsto \|[A, A^*]\|^2$ is the norm squared of a momentum map (Proposition 2.2), with set of critical points exactly equal to \mathcal{N}_d (Theorem 2.3), it follows by a more general result of Duistermaat (see the expository work of Lerman [40]) that gradient descent gives a strong deformation retract of $\mathbb{C}^{d \times d}$ onto \mathcal{N}_d . One can also deduce this from the work above: by Theorem 2.6, we have a well-defined function $\mathcal{F} : \mathbb{C}^{d \times d} \times [0, \infty] \rightarrow \mathcal{N}_d$ induced by gradient descent, which obviously fixes \mathcal{N}_d , and the arguments in [40] show that the map is continuous. Moreover, this restricts to a strong deformation retract $\mathcal{M}_d \times [0, \infty] \rightarrow \mathcal{N}_d \setminus \{0\}$, by Corollary 2.8. As $\mathcal{N}_d \setminus \{0\}$ is a cone over \mathcal{UN}_d , the former also strong deformation retracts onto the latter. Concatenating these two strong deformation retracts gives a strong deformation retract $\mathcal{M}_d \rightarrow \mathcal{UN}_d$. \square

In particular, \mathcal{UN}_d is homotopy equivalent to \mathcal{M}_d , so our goal of characterizing the topology of the former space reduces to understanding that of the latter space. From such an understanding, we will deduce the main theorem of this subsection, stated below. In the following, we use $\pi_k(\mathcal{X}, x_0)$ to denote the k th homotopy group of a space \mathcal{X} with respect to a basepoint $x_0 \in \mathcal{X}$, and write $\pi_k(\mathcal{X})$ in the case

that \mathcal{X} is path connected (in which case the result is independent of basepoint, up to isomorphism) – we refer the reader to [28, Chapter 4] for basic terminology and properties. We say that \mathcal{X} is k -connected if $\pi_k(\mathcal{X}, x_0)$ is the trivial group.

Theorem 3.9. *The space \mathcal{UN}_d is k -connected for all $k \leq 2d - 2$.*

Remark 3.10. In particular, \mathcal{UN}_d is connected for all d . Moreover, \mathcal{UN}_d is simply connected (i.e., $\pi_1(\mathcal{UN}_d)$ is also trivial) for all $d \geq 2$.

The proof will use two auxiliary topological results. The first follows from more general results on *nilpotent cones*, which are classical. We use [35] as a general reference and explain how to deduce this particular result from the general results therein.

Lemma 3.11. *The space \mathcal{P}_d of nilpotent matrices in $\mathbb{C}^{d \times d}$ is an irreducible variety of complex dimension $d(d - 1)$.*

Proof. We apply the general nilpotent cone theory to the Lie group of invertible matrices $\mathrm{GL}_d(\mathbb{C})$, in which case the nilpotent cone is exactly \mathcal{P}_d . Then the fact that \mathcal{P}_d is an irreducible variety is [35, Lemma 6.2]. By [35, Theorem 6.4], the dimension of \mathcal{P}_d is twice the dimension of the maximal unipotent subalgebra of the Lie algebra $\mathbb{C}^{d \times d}$, namely the subalgebra of strictly upper triangular matrices (i.e., with zeros on the diagonal). This subalgebra has complex dimension $1 + 2 + \cdots + (d - 1) = \frac{1}{2}d(d - 1)$. \square

The following is a standard application of transversality (see [39, Chapter 6] and [32, Chapter 3]). Special cases of the result appear in, for example, [23, Theorem 2.3] and [18, Theorem 1.1.4]. We give a proof sketch here for the sake of convenience.

Lemma 3.12. *Let \mathcal{X} be a connected smooth manifold and let $\mathcal{Y} \subset \mathcal{X}$ be a union of smooth submanifolds, $\mathcal{Y} = \mathcal{Y}_1 \cup \cdots \cup \mathcal{Y}_\ell$, such that each \mathcal{Y}_j has codimension greater than or equal to m in \mathcal{X} . Then $\pi_k(\mathcal{X} \setminus \mathcal{Y})$ is isomorphic to $\pi_k(\mathcal{X})$ for all $k \leq m - 2$.*

Proof. We will show that the inclusion map $\iota : \mathcal{X} \setminus \mathcal{Y} \hookrightarrow \mathcal{X}$ induces a bijection between homotopy groups.

To establish surjectivity, we will show that any map $f : S^k \rightarrow \mathcal{X}$ is homotopic to a map $S^k \rightarrow \mathcal{X} \setminus \mathcal{Y}$. To do so, we apply the Whitney Approximation Theorem [39, Theorem 6.26] to homotope f to a smooth map. By the version of the corollary of the Transversality Theorem given in [32, Theorem 2.5], together with the argument in the proof of the Transversality Homotopy Theorem [39, Theorem 6.36], the resulting map is then homotopic to a map $S^k \rightarrow \mathcal{X}$ which is transverse to each submanifold \mathcal{Y}_j . By the codimensionality constraint, this is only possible if the image of S^k is disjoint from each \mathcal{Y}_j . This shows that f is homotopic to a map whose image is disjoint from \mathcal{Y} .

Next, we show that the map induced by ι is injective. That is, if maps $f_0, f_1 : S^k \rightarrow \mathcal{X}$ are homotopic, and, without loss of generality (by the above), $f_0(S^k) \cap \mathcal{Y} = f_1(S^k) \cap \mathcal{Y} = \emptyset$, then they are homotopic in $\mathcal{X} \setminus \mathcal{Y}$. This is done by applying similar arguments to the above to the homotopy $f : S^k \times [0, 1] \rightarrow \mathcal{X}$; in particular, this map may be homotoped without destroying transversality at the boundary $S^k \times \{0, 1\}$ [32, Ch. 3, Theorem 2.1]. \square

Proof of Theorem 3.9. By Corollary 3.8, it suffices to show that \mathcal{M}_d is k -connected for all $2d - 2$. By a theorem of Whitney, the algebraic variety \mathcal{P}_d can be expressed as a disjoint union of smooth manifolds [59, Theorem 2], and, by Lemma 3.11, each of these has real codimension at least

$$\dim(\mathbb{C}^{d \times d}) - \dim(\mathcal{P}_d) = 2d^2 - 2d(d - 1) = 2d.$$

The theorem then follows from Lemma 3.12, since $\mathbb{C}^{d \times d}$ is k -connected for all k . \square

3.3. Topology of real unit norm normal matrices

Let $\mathcal{UN}_d^{\mathbb{R}}$ denote the space of real, normal $d \times d$ matrices with Frobenius norm equal to one (so $\mathcal{UN}_d^{\mathbb{R}} \subset \mathcal{UN}_d$). Adapting the arguments from the previous subsection, we will show the following.




Figure 5. Consider the space $\mathcal{U}_2^{\mathbb{R}}$ of 2×2 real matrices with Frobenius norm 1. Since $\mathcal{U}_2^{\mathbb{R}}$ is a copy of the 3-sphere, we can stereographically project to \mathbb{R}^3 . The image under this projection of the unit-norm nilpotent matrices is shown in blue, and the image of $\mathcal{UN}_2^{\mathbb{R}}$ is shown in pink. Specifically, the pink plane (which is the $y = z$ plane) is the image of the symmetric matrices and the pink loop is the image of the normal matrices of the form $\begin{bmatrix} a & b \\ -b & a \end{bmatrix}$.

Theorem 3.13. The space $\mathcal{UN}_d^{\mathbb{R}}$ is k -connected for all $k \leq d - 2$.

Remark 3.14. It follows from the theorem that $\mathcal{UN}_d^{\mathbb{R}}$ is path connected for $d \geq 2$ and simply connected for $d \geq 3$. These results are tight:

- $\mathcal{UN}_1^{\mathbb{R}} \approx \{\pm 1\}$ is not path connected.
- $\mathcal{UN}_2^{\mathbb{R}}$ is not simply connected. This is illustrated in Figure 5.

The proof of the theorem follows the same general steps as that of Theorem 3.9. Let $\mathcal{P}_d^{\mathbb{R}}$ denote the $d \times d$ real nilpotent matrices, and let $\mathcal{M}_d^{\mathbb{R}} = \mathbb{R}^{d \times d} \setminus \mathcal{P}_d^{\mathbb{R}}$ denote the set of non-nilpotent matrices. By the same arguments used in the previous subsection, $\mathcal{M}_d^{\mathbb{R}}$ deformation retracts onto $\mathcal{UN}_d^{\mathbb{R}}$, so it suffices to prove that $\mathcal{M}_d^{\mathbb{R}}$ is k -connected for all $k \leq d - 2$.

The main difference in the real case is that an analogue of Lemma 3.11 does not follow from general facts of nilpotent cones described in [35], as the results therein are valid over algebraically closed fields. We obtain a decomposition of $\mathcal{P}_d^{\mathbb{R}}$ in analogy with the Whitney decomposition used in the proof of Theorem 3.9 from results of [29] and [8].

Lemma 3.15. The set of nilpotent matrices $\mathcal{P}_d^{\mathbb{R}}$ is a union of smooth submanifolds of $\mathbb{R}^{d \times d}$, each of which has codimension at least d .

Proof. It follows from a general theory of real reductive Lie group actions developed in [29] that $\mathbb{R}^{d \times d} \setminus \{0\}$ decomposes as a union of $\mathrm{GL}_d(\mathbb{R})$ -invariant (with respect to the conjugation action) smooth submanifolds $S_0 \cup S_1 \cup \cdots \cup S_k$, where S_0 is exactly the open submanifold $\mathcal{M}_d^{\mathbb{R}}$ – see also [8, Section 1] for a short exposition of these ideas. It is shown in [8, Section 1.2] that (for the specific example of the conjugation action on $\mathbb{R}^{d \times d}$) the remaining submanifolds $S_i, i > 0$, are parameterized by Jordan canonical forms of nilpotent matrices. That is, fixing such a Jordan matrix J , we consider the corresponding set of nilpotent matrices as the homogeneous space $\mathrm{GL}_d(\mathbb{R})/\mathrm{stab}(J)$, where $\mathrm{stab}(J)$ is the stabilizer of J under the conjugation action. To complete the proof, it suffices to show that the dimension of such a homogeneous space is at most $d^2 - d$, that is, to show that the stabilizer of any such Jordan matrix is at least of dimension d .

Let us now establish the claim made above. A nilpotent Jordan matrix J necessarily has all zeros on its diagonal, and is therefore characterized by the pattern of ones in the super diagonal (i.e., by the size of its Jordan blocks). An invertible real matrix $A = (a_{ij})_{i,j=1}^d$ lies in the stabilizer of J if and only if $AJ = JA$. This matrix equation gives several constraints in the entries of A , and the number of independent constraints determines the dimension of the stabilizer.

In particular, since we aim to determine a lower bound on codimension, it suffices to consider the Jordan matrix which produces the largest number of constraints: that is, when J is the matrix whose superdiagonal consists of all ones (i.e., it has a single Jordan block). It is a standard fact (see, e.g., [24, Theorem 9.1.1]) that, for this J , solutions of the equation $AJ = JA$ must be upper triangular Toeplitz matrices. In other words, elements of $\text{stab}(J)$ are of the form

$$\begin{bmatrix} a_1 & a_2 & a_3 & \cdots & a_{d-1} & a_d \\ 0 & a_1 & a_2 & \cdots & a_{d-2} & a_{d-1} \\ 0 & 0 & a_1 & \cdots & a_{d-3} & a_{d-2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & a_1 & a_2 \\ 0 & 0 & 0 & \cdots & 0 & a_1 \end{bmatrix}.$$

Clearly, then, $\dim(\text{stab}(J)) = d$. This implies that the codimension of the associated submanifold is d . Since this is the submanifold of smallest codimension, this completes the proof. \square

Proof of Theorem 3.13. By the discussion above, it suffices to prove that $\mathcal{M}_d^{\mathbb{R}}$ is k -connected for all $k \leq d - 2$. In light of Lemma 3.15, the same transversality argument as was used in the proof of Theorem 3.9 can be used here. \square

4. Balanced matrices and weighted digraphs

As was described in the introduction, the techniques and results that we have developed for normal matrices can be adapted to the setting of weighted digraphs. The naturality of such an application follows from the following observation. Notice that the diagonal entries of $\mu(A) = [A, A^*]$ are of the form $\|A_i\|^2 - \|A^i\|^2$, where A_i is the i th row of A and A^i is the i th column. Hence, if $A \in \mathbb{C}^{d \times d}$ is normal, then $\|A_i\|^2 = \|A^i\|^2$ for all $i = 1, \dots, d$. This suggests a certain *balancing condition*, as we expand on below.

Suppose that \mathcal{G} is a weighted, directed graph and \widehat{A} is its associated adjacency matrix; that is, the (i, j) entry of \widehat{A} is the (non-negative) weight of the directed edge from vertex i to vertex j if such an edge exists, and zero if there is no such edge. In particular, the entries of \widehat{A} are non-negative real numbers. If A is the matrix whose entries are the square roots of the entries of \widehat{A} , then $\|A_i\|^2 = \|A^i\|^2$ says that the i th vertex v_i of \mathcal{G} is *balanced*: the sum of the weights of the edges coming into v_i equals the sum of the weights of the edges leaving v_i . In other words, every real normal matrix A corresponds to a balanced, weighted, directed (multi-)graph⁵ \mathcal{G} by interpreting the component-wise square of A as the adjacency matrix of \mathcal{G} . Moreover, the gradient descent procedures described in the previous sections give ways of balancing a given weighted, directed graph.

However, balancing a graph by gradient descent of E or \overline{E} has some undesirable features. First, the condition that A is normal is stronger than the condition that \mathcal{G} is balanced;⁶ second, the gradient flow is not guaranteed to ensure that a zero entry in the adjacency matrix will stay zero, so the limiting balanced graph may have sprouted new edges (and even loop edges) not present in the initial graph.

⁵If A has nonzero entries on its diagonal, the resulting graph will have loop edges at the corresponding vertices.

⁶For example, in the case when all weights are 1, normality of A implies that every pair of vertices (not necessarily distinct) has the same number of common out-neighbors as common in-neighbors.

For applications to balancing graphs, then, the natural energy to consider is not the non-normal energy E , but rather the *unbalanced energy* $B : \mathbb{C}^{d \times d} \rightarrow \mathbb{R}$ defined by

$$B(A) = \|\operatorname{diag}(AA^* - A^*A)\|^2 = \sum_{i=1}^d \left(\|A_i\|^2 - \|A^i\|^2 \right)^2,$$

where we use A_i for the i th row of A and A^i for the i th column. We will say that A is *balanced* if $B(A) = 0$.

We now describe this function from the perspective of symplectic geometry and GIT. Following a general theme of the paper, these observations are not really essential in what follows, but they provided inspiration, especially in light of Kirwan's fundamental work [37]. Let $\operatorname{DSU}(d)$ be the subgroup of $\operatorname{SU}(d)$ consisting of diagonal matrices. Then $\operatorname{DSU}(d) \approx \operatorname{U}(1)^{d-1}$ is the standard maximal torus of $\operatorname{SU}(d)$. The restriction of the conjugation action of $\operatorname{SU}(d)$ on $\mathbb{C}^{d \times d}$ gives a Hamiltonian action of $\operatorname{DSU}(d)$ on $\mathbb{C}^{d \times d}$ with momentum map $\mu_\Delta : \mathbb{C}^{d \times d} \rightarrow \mathbb{R}^d$ given by composing the momentum map μ of the $\operatorname{SU}(d)$ action with orthogonal projection to $\mathfrak{su}(d)^* \subset \mathfrak{su}(d)^*$ (see, e.g., [4, Proposition III.1.10]). Under the identification of $\mathfrak{su}(d)^*$ with the traceless Hermitian matrices, $\mathfrak{su}(d)^*$ corresponds to the traceless, diagonal, real matrices, so we have

$$\mu_\Delta(A) = \operatorname{diag}(\mu(A)) = \operatorname{diag}([A, A^*])$$

and $B(A) = \|\mu_\Delta(A)\|^2$. The GIT version of the foregoing is that the diagonal subgroup $\operatorname{DSL}_d(\mathbb{C}) \subset \operatorname{SL}_d(\mathbb{C})$ has an algebraic action by conjugation on $\mathbb{C}^{d \times d}$ (or, in Section 4.2, on $\mathbb{P}(\mathbb{C}^{d \times d})$).

4.1. Balancing matrices by gradient descent

As in the case of E , all critical points of B are global minima:

Theorem 4.1. *The only critical points of B are the global minima; that is, the balanced matrices.*

Proof. We first show that the gradient of the balanced energy is given by

$$\nabla B(A) = -4[A, \operatorname{diag}([A, A^*])]. \quad (4.1)$$

We write $B = N \circ \operatorname{diag} \circ \mu$, where μ is the momentum map (2.1), we consider the diagonalization operator as a linear map $\operatorname{diag} : \mathbb{C}^{d \times d} \rightarrow \mathbb{C}^{d \times d}$, and N is the norm-squared map, as in the proof of Theorem 2.3. Following the logic of that proof, we then have that

$$\nabla B(A) = D\mu(A)^\vee \operatorname{diag}^\vee \nabla N(\operatorname{diag} \circ \mu(A)),$$

where the superscripts once again denote adjoints with respect to $\langle \cdot, \cdot \rangle$. It is not hard to show that the map diag is self-adjoint and idempotent. Then

$$\begin{aligned} \nabla B(A) &= D\mu(A)^\vee \operatorname{diag}(2 \cdot \operatorname{diag} \circ \mu(A)) = 2D\mu(A)^\vee \operatorname{diag}(\mu(A)) \\ &= 2[\operatorname{diag}(\mu(A)) + \operatorname{diag}(\mu(A))^*, A] = -4[A, \operatorname{diag}([A, A^*])]. \end{aligned}$$

The above shows that we have a critical point of B exactly when

$$0 = [A, \operatorname{diag}([A, A^*])].$$

Since the entries of $[A, \operatorname{diag}([A, A^*])]$ are of the form

$$a_{ij} \left((\|A_i\|^2 - \|A^i\|^2) - (\|A_j\|^2 - \|A^j\|^2) \right), \quad (4.2)$$

this means that $\|A_i\|^2 - \|A^i\|^2 = \|A_j\|^2 - \|A^j\|^2$ for all i and j such that $a_{ij} \neq 0$.

In other words, A is a critical point of B if and only if $\|A_i\|^2 - \|A^i\|^2$ is independent of i . However, since


$$\sum_{i=1}^d (\|A_i\|^2 - \|A^i\|^2) = \sum_{i=1}^d \|A_i\|^2 - \sum_{i=1}^d \|A^i\|^2 = \|A\|^2 - \|A\|^2 = 0,$$

this can only happen if all $\|A_i\|^2 - \|A^i\|^2 = 0$; that is, if A is balanced. \square

Remark 4.2. Theorem 4.1 shows that B is an invex function, but it is not quasiconvex. To see this, consider the matrices

$$A_0 = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix} \quad \text{and} \quad A_1 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix},$$

which are the (entrywise square roots of the) adjacency matrices of the balanced graphs



respectively. Then $B(A_0) = 0 = B(A_1)$, but $B((1-t)A_0 + tA_1) = 8t^2(1-t)^2 > 0$ for all $0 < t < 1$.

As in the case of E , we can find global minima of B by gradient descent. Specifically, let $\mathcal{F} : \mathbb{C}^{d \times d} \times [0, \infty) \rightarrow \mathbb{C}^{d \times d}$ be the negative gradient flow of B :

$$\mathcal{F}(A_0, 0) = A_0 \quad \frac{d}{dt} \mathcal{F}(A_0, t) = -\nabla B(\mathcal{F}(A_0, t)).$$

Since B is a real polynomial function on $\mathbb{C}^{d \times d}$, Theorem 4.1 implies that $\lim_{t \rightarrow \infty} \mathcal{F}(A_0, t)$ is always well-defined and normal. Since the real matrices stay real under gradient flow, this limit will be real whenever A_0 is.

Moreover,

$$\nabla B(A) = -4[A, \text{diag}([A, A^*])] = 4 \left(\frac{d}{d\epsilon} \Big|_{\epsilon=0} e^{\epsilon \text{diag}([A, A^*])} \cdot A \right) \quad (4.3)$$

is tangent to the orbit of the diagonal subgroup $\text{DSL}_d(\mathbb{C}) \leq \text{SL}_d(\mathbb{C})$ acting by conjugation on $\mathbb{C}^{d \times d}$. In particular, flowing A_0 by the gradient flow of B preserves not just the eigenvalues of A_0 , but also all principal minors of A_0 , including the diagonal entries of A_0 .

From the expression (4.2) for the entries of $-\frac{1}{4} \nabla B(A)$ we see that, if there is $t_0 \geq 0$ so that the (i, j) entry in $\mathcal{F}(A_0, t_0)$ vanishes, then the (i, j) entry of $\mathcal{F}(A_0, t)$ will vanish for all $t \geq t_0$. In graph terms, the gradient flow of B cannot sprout new edges in the graph. This also means that if A_0 is real, its entries cannot change sign under gradient descent of B . Thus, we have proved:

Theorem 4.3. For any $A_0 \in \mathbb{C}^{d \times d}$, the matrix $A_\infty := \lim_{t \rightarrow \infty} \mathcal{F}(A_0, t)$ exists, is balanced, has the same eigenvalues and principal minors as A_0 , and has zero entries wherever A_0 does. If A_0 is real, then so is A_∞ , and if A_0 has all non-negative entries, then so does A_∞ .

When we take A_0 to be the entrywise square root of the adjacency matrix of some weighted, directed graph $\mathcal{G}_0 = (\mathcal{V}_0, \mathcal{E}_0, w_0)$, then we can sensibly interpret $A_\infty = \lim_{t \rightarrow \infty} \mathcal{F}(A_0, t)$ as the entrywise square root

of the adjacency matrix of some balanced, weighted, directed graph $\mathcal{G}_\infty = (\mathcal{V}_\infty, \mathcal{E}_\infty, w_\infty)$ with $\mathcal{V}_\infty = \mathcal{V}_0$ and $\mathcal{E}_\infty \subseteq \mathcal{E}_0$. In other words, gradient descent of B balances \mathcal{G}_0 without introducing any new edges.

Remark 4.4. An important consideration in the applied literature on graph balancing is that algorithms are *local*, in the sense that iterative updates are only performed based on node-level information [51, 33, 26]. This is due both to practical constraints on data acquisition, as well as the need for parallelizability in computation. Observe from the structure of the gradient of the unbalanced energy that the gradient descent approach to graph balancing is not local in the sense described above, but is *semilocal* in the sense that updates only depend on edge-level information. While this paper is concerned with theory and makes no claims to efficiency or practicality of the algorithm, the useful properties of the gradient flow of B suggest that it may be interesting to explore its viability in real-world applications.

In the case of gradient descent of E , we saw that all nilpotent matrices flowed to the zero matrix. We see the same phenomenon here: if \mathcal{G}_0 is a weighted, directed, acyclic graph (DAG), then its adjacency matrix is nilpotent, as is the entrywise square root A_0 . The gradient flow $\mathcal{F}(A_0, t)$ will limit to the zero matrix, which makes sense: the only way to balance a weighted DAG is by driving all the weights to zero.

4.2. Preserving weights

Weighted DAGs provide an extreme example of the general phenomenon that gradient descent of B decreases the Frobenius norm. In graph terms, if A_0 is the entrywise square root of the adjacency matrix of a weighted, directed graph \mathcal{G}_0 , then the squared Frobenius norm

$$\|A_0\|^2 = \sum_{i,j} |a_{ij}|^2 = \sum_{i,j} a_{ij}^2$$

is precisely the sum of the weights in \mathcal{G}_0 . If the weights correspond to, for example, mass traversing between nodes in a network, then it may not make sense to balance the flows in the network by reducing the total mass in the system.

In order to preserve the sum of weights on \mathcal{G}_0 , we consider $\bar{B} : \mathcal{U}_d \rightarrow \mathbb{R}$, the restriction of B to \mathcal{U}_d , and its gradient descent $\bar{\mathcal{F}} : \mathcal{U}_d \times [0, \infty) \rightarrow \mathcal{U}_d$ given by

$$\bar{\mathcal{F}}(A_0, 0) = A_0 \quad \frac{d}{dt} \bar{\mathcal{F}}(A_0, t) = -\text{grad } \bar{B}(\bar{\mathcal{F}}(A_0, t)).$$

Theorem 4.5. *For any non-nilpotent $A_0 \in \mathcal{U}_d$, the matrix $A_\infty := \lim_{t \rightarrow \infty} \bar{\mathcal{F}}(A_0, t)$ exists, is balanced, has Frobenius norm 1, and has zero entries wherever A_0 does. If A_0 is real, so is A_∞ , and if A_0 has all non-negative entries, then so does A_∞ .*

In graph terms, if A_0 is the entrywise square root of an adjacency matrix for \mathcal{G}_0 with total weight 1, then A_∞ is the entrywise square root of the adjacency matrix for a balanced graph \mathcal{G}_∞ with total weight 1 whose vertices are the same as the vertices of \mathcal{G}_0 and whose edges are a subset of the edges of \mathcal{G}_0 . That is, gradient descent of \bar{B} balances \mathcal{G}_0 without introducing any new edges and without losing any overall weight.

The strategy for proving Theorem 4.5 is the same as for Theorem 3.1. The existence of a unique limit point A_∞ follows from the fact that \bar{B} is a polynomial function on \mathcal{U}_d , and hence has a Łojasiewicz exponent. The bulk of the argument is in showing that the gradient flow preserves non-nilpotency and that the nonminimizing critical points are nilpotent. The rest of the theorem will follow from the structure of $\text{grad } \bar{B}$ and the fact that the real submanifold of \mathcal{U}_d is invariant under gradient flow.

First, we compute the intrinsic gradient of \bar{B} , which follows the same pattern as $\text{grad } E$:

Proposition 4.6. *The intrinsic gradient of \bar{B} on \mathcal{U}_d is*

$$\text{grad } \bar{B}(A) = -4([A, \text{diag}([A, A^*])] + \bar{B}(A)A).$$

Proof. We know that

$$\text{grad } \overline{\mathbf{B}}(A) = \nabla \mathbf{B}(A) - \langle \nabla \mathbf{B}(A), A \rangle A,$$

so the key is to use (4.1) and the fact that the diagonal of $[A, A^*]$ is real to compute

$$\begin{aligned} \langle \nabla \mathbf{B}(A), A \rangle &= -4\text{Re tr}([A, \text{diag}([A, A^*])]^* A) \\ &= -4\text{Re tr}(\text{diag}([A, A^*]) A^* A - A^* \text{diag}([A, A^*]) A) \\ &= 4\text{Re tr}(\text{diag}([A, A^*]) [A, A^*]) \\ &= 4\text{Re tr}(\text{diag}([A, A^*]) \text{diag}([A, A^*])) \\ &= 4\|\text{diag}([A, A^*])\|^2 \\ &= 4\overline{\mathbf{B}}(A) \end{aligned}$$

using the linearity and cyclic invariance of trace. \square

Each entry of $\text{grad } \overline{\mathbf{B}}(A)$ is a scalar multiple of the corresponding entry of A , so the fact that the negative gradient flow $\overline{\mathcal{F}}$ preserves zero entries and cannot change the sign of real entries follows immediately.

Next, we prove an analog of Lemma 3.5. Recall that $s(A) = \sum |\lambda_i|^2$ is the sum of the squares of the absolute values of the eigenvalues of A .

Lemma 4.7. *For any $A \in \mathcal{U}_d$,*

$$\langle -\text{grad } \overline{\mathbf{B}}(A), \text{grad } s(A) \rangle = 8s(A)\overline{\mathbf{B}}(A).$$

Proof. The proof exactly parallels the proof of Lemma 3.5 by substituting \mathbf{B} , $\overline{\mathbf{B}}$, and (4.3) for \mathbf{E} , $\overline{\mathbf{E}}$, and (2.6), respectively. \square

Since $\langle -\text{grad } \overline{\mathbf{B}}(A), \text{grad } s(A) \rangle = 8s(A)\overline{\mathbf{B}}(A) \geq 0$, $s(A)$ must be nondecreasing along the negative gradient flow lines of $\overline{\mathbf{B}}$, so we have proved:

Proposition 4.8. *If $A_0 \in \mathcal{U}_d$ is non-nilpotent, then so is $A_t := \overline{\mathcal{F}}(A_0, t)$ and so is $A_\infty := \lim_{t \rightarrow \infty} \overline{\mathcal{F}}(A_0, t)$.*

We know the balanced matrices are exactly the global minima of $\overline{\mathbf{B}}$. Proposition 4.6 implies that A is a critical point of $\overline{\mathbf{B}}$ if and only if

$$0 = [A, \text{diag}([A, A^*])] + \overline{\mathbf{B}}(A)A.$$

When A is a non-minimizing critical point, $\overline{\mathbf{B}}(A) \neq 0$ and the same Jacobson's Lemma argument as in Proposition 3.7 shows that A is nilpotent, proving:

Proposition 4.9. *All non-minimizing critical points of $\overline{\mathbf{B}}$ are nilpotent.*

This completes the proof of Theorem 4.5.

Figure 1 shows an application of this approach to balancing graphs, and Figure 6 shows a much larger example. In both cases, up to an overall normalization to ensure $\|A_0\| = 1$, the non-zero entries in the starting matrix A_0 were populated by the absolute values of standard Gaussians.

4.3. Topology of unit norm balanced graphs

Let \mathcal{UB}_d denote the space of balanced $d \times d$ matrices of unit Frobenius norm, and let $\mathcal{UB}_d^{\mathbb{R}}$ denote the subspace of balanced matrices with real entries. The topology of these spaces is tied to the topology of the relevant spaces of normal matrices, as we record in the following theorem.




Figure 6. Balancing a larger graph by the flow $\overline{\mathcal{F}}$, with A_0 on the left and $A_\infty = \lim_{t \rightarrow \infty} \overline{\mathcal{F}}(A_0, t)$ on the right. The thickness of each edge is proportional to its weight. The underlying graph is a random planar graph with 100 vertices and 284 edges, constructed as the 1-skeleton of the Delaunay triangulation of 100 random points in the square; to make the visualization more comprehensible, the graph that is shown is a spring embedding, so the vertices are not at the locations of the original random points in the square.

Theorem 4.10. The spaces \mathcal{UN}_d and \mathcal{UB}_d are homotopy equivalent. Similarly, the spaces $\mathcal{UN}_d^{\mathbb{R}}$ and $\mathcal{UB}_d^{\mathbb{R}}$ are homotopy equivalent.

Proof. By Theorem 3.1 and Theorem 4.5, respectively, \mathcal{UN}_d and \mathcal{UB}_d are both deformation retracts of the space of non-nilpotent unit norm matrices. The same theorems give the result in the real case. \square

Acknowledgments. We are grateful to Malbor Asllani for catalyzing this line of inquiry, to Chris Peterson for enlightening conversations, and to the anonymous referees for their careful reading and thoughtful comments which have made this a better paper. We would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the program *New equivariant methods in algebraic and differential geometry*, where some of the work on this paper was undertaken.

Competing interest. The authors have no competing interests to declare.

Funding statement. This work was supported by EPSRC grant EP/R014604/1 and by grants from the National Science Foundation (DMS–2107808 and DMS–2324962, Tom Needham; DMS–2107700, Clayton Shonkwiler).

References

- [1] P.-A. Absil and K. Kurdyka, ‘On the stable equilibrium points of gradient systems’, *Syst. Control Lett.* **55**(7) (2006), 573–577.
- [2] M. Asllani and T. Carletti, ‘Topological resilience in non-normal networked systems’, *Phys. Rev. E* **97**(4) (2018), 042302.
- [3] M. Asllani, R. Lambiotte and T. Carletti, ‘Structure and dynamical behavior of non-normal networks’, *Sci. Adv.* **4**(12) (2018), eaau9403.
- [4] M. Audin, *Torus Actions on Symplectic Manifolds* (Birkhäuser, Basel, 2nd rev. ed., 2004).
- [5] F. L. Bauer and C. T. Fike, ‘Norms and exclusion theorems’, *Numer. Math.* **2**(1) (1960), 137–141.
- [6] A. Ben-Israel and B. Mond, ‘What is invexity?’, *J. Austral. Math. Soc. Ser. B Appl. Math.* **28**(1) (1986), 1–9.
- [7] B. G. Bodmann and J. Haas, ‘Frame potentials and the geometry of frames’, *J. Fourier Anal. Appl.* **21**(6) (2015), 1344–1383.
- [8] C. Böhm and R. A. Lafuente, ‘Real geometric invariant theory’, in O. Dearnicott, W. Tuschmann, Y. Nikolayevsky, T. Leistner and D. Crowley (eds.), *Differential Geometry in the Large*, No. 463 of London Mathematical Society Lecture Note Series (Cambridge Univ. Press, Cambridge, 2021), 11–49.
- [9] J. Cahill, D. G. Mixon and N. Strawn, ‘Connectivity and irreducibility of algebraic varieties of finite unit norm tight frames’, *SIAM J. Appl. Algebra Geom.* **1**(1) (2017), 38–72.

- [10] M. Calvo, A. Iserles and A. Zanna, 'Numerical solution of isospectral flows', *Math. Comp.* **66**(220) (1997), 1461–1486.
- [11] M. T. Chu, 'Least squares approximation by real normal matrices with specified spectrum', *SIAM J. Matrix Anal. Appl.* **12**(1) (1991), 115–127.
- [12] M. T. Chu, 'Linear algebra algorithms as dynamical systems', *Acta Numer.* **17** (2008), 1–86.
- [13] B. D. Craven, 'Duality for generalized convex fractional programs', in S. Schaible and W. T. Ziemba (eds.), *Generalized Concavity in Optimization and Economics: Proc. NATO Adv. Study Inst., Univ. British Columbia, Vancouver, Canada, Aug. 4–15, 1980* (Academic Press, New York, 1981), 473–489.
- [14] B. D. Craven and B. M. Glover, 'Invex functions and duality', *J. Austral. Math. Soc. Ser. A Pure Math. Stat.* **39**(1) (1985), 1–20.
- [15] R. W. Daniel and B. Kouvaritakis, 'The choice and use of normal approximations to transfer-function matrices of multivariable control systems', *Int. J. Control* **37**(5) (1983), 1121–1133.
- [16] R. W. Daniel and B. Kouvaritakis, 'Analysis and design of linear multivariable feedback systems in the presence of additive perturbations', *Int. J. Control* **39**(3) (1984), 551–580.
- [17] P. Deift, T. Nanda and C. Tomei, 'Ordinary differential equations and the symmetric eigenvalue problem', *SIAM J. Numer. Anal.* **20**(1) (1983), 1–22.
- [18] J. Ebert, 'A lecture course on cobordism theory', unpublished lecture notes (2012). URL: https://ivv5hpp.uni-muenster.de/u/jeber_02/skripten/bordism-skript.pdf.
- [19] L. Elsner and K. D. Ikramov, 'Normal matrices: an update', *Linear Algebra Appl.* **285**(1–3) (1998), 291–303.
- [20] J. M. Fisher, 'Morse theory with the norm-square of a hyperKähler moment map', *Q. J. Math.* **65**(1) (2014), 149–173.
- [21] S. Friedland, 'Normal matrices and the completion problem', *SIAM J. Matrix Anal. Appl.* **23**(3) (2002), 896–902.
- [22] R. Gabriel, 'The normal ΔH -matrices with connection to some Jacobi-like methods', *Linear Algebra Appl.* **91** (1987), 181–194.
- [23] C. Godbillon, *Éléments de topologie algébrique* (Hermann, Paris, 1971).
- [24] I. Gohberg, P. Lancaster and L. Rodman, *Invariant Subspaces of Matrices with Applications*, No. 51 of Classics in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM, Philadelphia, PA, 2006).
- [25] N. Guglielmi and C. Scalone, 'Computing the closest real normal matrix and normal completion', *Adv. Comput. Math.* **45** (2019), 2867–2891.
- [26] C. N. Hadjicostis and A. Rikos, 'Distributed strategies for balancing a weighted digraph', in *Proc. 20th Mediterr. Conf. Control Autom. (MED)* (IEEE, 2012), 1141–1146.
- [27] M. A. Hanson, 'On sufficiency of the Kuhn–Tucker conditions', *J. Math. Anal. Appl.* **80**(2) (1981), 545–550.
- [28] A. Hatcher, *Algebraic Topology* (Cambridge Univ. Press, Cambridge, 2002).
- [29] P. Heinzner, G. W. Schwarz and H. Stötzl, 'Stratifications with respect to actions of real reductive groups', *Compos. Math.* **144**(1) (2008), 163–185.
- [30] P. Henrici, 'Bounds for iterates, inverses, spectral variation and fields of values of non-normal matrices', *Numer. Math.* **4**(1) (1962), 24–40.
- [31] N. J. Higham, 'Matrix nearness problems and applications', in M. J. C. Gover and S. Barnett (eds.), *Applications of Matrix Theory: Proc. Conf. Univ. Bradford, July 1988*, vol. 22 of *IMA Conf. Ser. New Ser.* (Clarendon Press, Oxford Univ. Press, New York, 1989), 1–27.
- [32] M. W. Hirsch, *Differential Topology*, No. 33 of Graduate Texts in Mathematics (Springer, New York, 2012).
- [33] H.-T. Loh, 'On a class of directed graphs—with an application to traffic-flow problems', *Oper. Res.* **18**(1) (1970), 87–94.
- [34] N. Jacobson, 'Rational methods in the theory of Lie algebras', *Ann. Math. (2)* **36**(4) (1935), 875–881.
- [35] J. C. Jantzen, 'Nilpotent orbits in representation theory', in J.-P. Anker and B. Orsted (eds.), *Lie Theory: Lie Algebras and Representations* (Birkhäuser, Boston, MA, 2004), 1–211.
- [36] I. Kaplansky, 'Jacobson's Lemma revisited', *J. Algebra* **62**(2) (1980), 473–476.
- [37] F. Kirwan, *Cohomology of Quotients in Symplectic and Algebraic Geometry*, vol. 31 of *Math. Notes* (Princeton Univ. Press, Princeton, NJ, 1984).
- [38] B. Kostant, 'Lie group representations on polynomial rings', *Amer. J. Math.* **85**(3) (1963), 327–404.
- [39] J. M. Lee, *Introduction to Smooth Manifolds*, No. 218 of Graduate Texts in Mathematics (Springer, New York, 2nd ed., 2013).
- [40] E. Lerman, 'Gradient flow of the norm squared of a moment map', *Enseign. Math.* **51** (2005), 117–127.
- [41] S. Łojasiewicz, 'Sur les trajectoires du gradient d'une fonction analytique', in *Seminari di Geometria 1982–1983* (Dip. Mat., Univ. Bologna, 1984), 115–117.
- [42] M. Marcus and H. Minc, *A Survey of Matrix Theory and Matrix Inequalities*, *Series in Advanced Mathematics* (Allyn and Bacon, Boston, MA, 1964).
- [43] D. G. Mixon, T. Needham, C. Shonkwiler and S. Villar, 'Three proofs of the Benedetto–Fickus theorem', in S. D. Casey, M. M. Dodson, P. J. S. G. Ferreira and A. Zayed (eds.), *Sampling, Approximation, and Signal Analysis: Harmonic Analysis in the Spirit of J. Rowland Higgins*, Applied and Numerical Harmonic Analysis (Birkhäuser, Cham, 2023), 371–391.
- [44] D. Mumford, J. Fogarty and F. Kirwan, *Geometric Invariant Theory*, vol. 34 of *Ergebnisse der Mathematik und ihrer Grenzgebiete* (Springer-Verlag, Berlin, 1994).
- [45] R. Muolo, M. Asllani, D. Fanelli, P. K. Maini and T. Carletti, 'Patterns of non-normality in networked systems', *J. Theoret. Biol.* **480** (2019), 81–91.
- [46] T. Needham and C. Shonkwiler, 'Symplectic geometry and connectivity of spaces of frames', *Adv. Comput. Math.* **47**(1) (2021), 5.

- [47] T. Needham and C. Shonkwiler, ‘Toric symplectic geometry and full spark frames’, *Appl. Comput. Harmon. Anal.* **61** (2022), 254–287.
- [48] A. Neeman, ‘The topology of quotient varieties’, *Ann. Math. (2)* **122**(2) (1985), 419–459.
- [49] L. Ness, ‘A stratification of the null cone via the moment map’, *Amer. J. Math.* **106**(6) (1984), 1281–1329.
- [50] S. Nosschese and L. Reichel, ‘The structured distance to normality of banded Toeplitz matrices’, *BIT Numer. Math.* **49** (2009), 629–640.
- [51] A. I. Rikos, T. Charalambous and C. N. Hadjicostis, ‘Distributed weight balancing over digraphs’, *IEEE Trans. Control Netw. Syst.* **1**(2) (2014), 190–201.
- [52] A. Ruhe, ‘Closest normal matrix finally found!’, *BIT Numer. Math.* **27** (1987), 585–598.
- [53] H. Rutishauser, ‘Solution of eigenvalue problems with the LR-transformation’, in *Further Contributions to the Solution of Simultaneous Linear Equations and the Determination of Eigenvalues*, vol. 49 of National Bureau of Standards Applied Mathematics Series (U.S. Govt. Printing Office, Washington, DC, 1958), 47–81.
- [54] C. Shonkwiler, *Geometry of Normal Matrices*. URL: <https://github.com/shonkwiler/normal-matrices-computations>.
- [55] R. P. Thomas, ‘Notes on GIT and symplectic reduction for bundles and varieties’, *Surv. Differ. Geom.* **10**(1) (2005), 221–273.
- [56] C. Tomei, ‘The Toda lattice, old and new’, *J. Geom. Mech.* **5**(4) (2013), 511–530.
- [57] D. S. Watkins, ‘Isospectral flows’, *SIAM Rev.* **26**(3) (1984), 379–391.
- [58] D. S. Watkins and L. Elsner, ‘On Rutishauser’s approach to self-similar flows’, *SIAM J. Matrix Anal. Appl.* **11**(2) (1990), 301–311.
- [59] H. Whitney, ‘Elementary structure of real algebraic varieties’, *Ann. Math.* **66**(3) (1957), 545–556.
- [60] C. T. Woodward, ‘The Yang–Mills heat flow on the moduli space of framed bundles on a surface’, preprint (2002). [arXiv:math/0211231](https://arxiv.org/abs/math/0211231). This is a preprint version of ‘The Yang–Mills heat flow on the moduli space of framed bundles on a surface’, *Amer. J. Math.* **128**(2) (2006), 311–369.