



Global Convergence of Hessenberg Shifted QR I: Exact Arithmetic

Jess Banks¹ · Jorge Garza-Vargas¹ · Nikhil Srivastava¹

Received: 12 July 2022 / Revised: 14 October 2023 / Accepted: 23 October 2023
© SFoCM 2024

Abstract

Rapid convergence of the shifted QR algorithm on symmetric matrices was shown more than 50 years ago. Since then, despite significant interest and its practical relevance, an understanding of the dynamics and convergence properties of the shifted QR algorithm on nonsymmetric matrices has remained elusive. We introduce a new family of shifting strategies for the Hessenberg shifted QR algorithm. We prove that when the input is a diagonalizable Hessenberg matrix H of bounded *eigenvector condition number* $\kappa_V(H)$ —defined as the minimum condition number of V over all diagonalizations $V D V^{-1}$ of H —then the shifted QR algorithm with a certain strategy from our family is guaranteed to converge rapidly to a Hessenberg matrix with a zero sub-diagonal entry, in exact arithmetic. Our convergence result is nonasymptotic, showing that the geometric mean of certain subdiagonal entries of H decays by a fixed constant in every QR iteration. The arithmetic cost of implementing each iteration of our strategy scales roughly logarithmically in the eigenvector condition number $\kappa_V(H)$, which is a measure of the nonnormality of H . The key ideas in the design and analysis of our strategy are: (1) we are able to precisely characterize when a certain shifting strategy based on Ritz values stagnates. We use this information to design certain “exceptional shifts” which are guaranteed to escape stagnation whenever it occurs. (2) We use higher degree shifts (of degree roughly $\log \kappa_V(H)$) to dampen transient effects

Communicated by Nick Higham.

Jess Banks: Supported by NSF GRFP Grant DGE-1752814 and NSF Grant CCF-2009011. Jorge Garza-Vargas: Supported by NSF Grant CCF-2009011. Nikhil Srivastava: Supported by NSF Grant CCF-2009011.

✉ Nikhil Srivastava
nikhil@math.berkeley.edu

Jess Banks
jess.m.banks@berkeley.edu

Jorge Garza-Vargas
jgarzavargas@berkeley.edu

¹ UC Berkeley, Berkeley, CA, USA

due to nonnormality, allowing us to treat nonnormal matrices in a manner similar to normal matrices.

Keywords Linear algebra · Numerical analysis · Dynamical systems · Nonnormality

Mathematics Subject Classification 65F15

Contents

- 1 Introduction
- 1.1 Statement of Results and Key Ideas
- 1.2 History and Related Work
- 2 Preliminaries and Notation
- 3 Normal Matrices
- 4 Nonnormal Matrices
- 4.1 Promising Ritz Values and Almost Monotonicity of the Potential
- 4.2 The Shifting Strategy
- 4.3 Efficiently Finding a Promising Ritz Value
- 4.4 Analysis of the Exceptional Shift
- References

1 Introduction

The Hessenberg shifted QR algorithm, discovered in the late 1950’s by Francis [27, 28] (see also Kublanovskaya [36]), is and has been for several decades the most widely used method for approximately¹ computing all of the eigenvalues of a dense matrix. It is implemented in all of the major software packages for numerical linear algebra and was listed as one of the “Top 10 algorithms of the twentieth century” [24, 45]. The algorithm is specified by a *shifting strategy*, which is an efficiently computable² function

$$\text{Sh} : \mathbb{H}^{n \times n} \rightarrow \mathcal{P}_k,$$

where $\mathbb{H}^{n \times n}$ is the set of $n \times n$ complex Hessenberg³ matrices and \mathcal{P}_k is the set of monic complex univariate polynomials of degree k , for some $k = k(n)$ typically much smaller than n . The word “shift” comes from the fact that when $k = 1$ we have $p_t(H_t) = H_t - s_t I$ ⁴ for some $s_t \in \mathbb{C}$. The algorithm then consists of the following discrete-time isospectral nonlinear dynamical system on $\mathbb{H}^{n \times n}$, given an initial condition H_0 :

¹ In the sense of backward error, i.e., exactly computing the eigenvalues of a nearby matrix.
² In this paper, we assume exact arithmetic with complex numbers and count arithmetic operations as a measure of complexity.
³ A matrix H is (upper) Hessenberg if $H(i, j) = 0$ whenever $i > j + 1$. Such matrices are “almost” upper triangular.
⁴ Hereon, following a convention from operator theory, we will use scalars to denote scalar multiples of the identity.

$$\begin{aligned} Q_t R_t &= p_t(H_t) & \text{where } p_t &= \text{Sh}(H_t), \\ H_{t+1} &= Q_t^* H_t Q_t, & t &= 0, 1, 2, \dots \end{aligned} \quad (1)$$

The first step in (1) is a QR decomposition, i.e. Q_t is unitary and R_t is upper triangular. It is not hard to see that this iteration preserves Hessenberg structure. Moreover, under the specification that the diagonal entries of R_t are positive, the QR decomposition is uniquely determined whenever $p_t(H_t)$ is invertible. Therefore, for the sake of simplicity, in this introduction we will ignore the case when $p_t(H_t)$ is singular (see Sect. 1.2 for a discussion).

The relevance of this iteration to the eigenvalue problem stems from two facts. First, every matrix $A \in \mathbb{C}^{n \times n}$ is unitarily similar to a Hessenberg matrix H_0 , and in exact arithmetic such a similarity can be computed exactly in $O(n^3)$ operations. Second, it was shown in [27, 36] that for the trivial “unshifted” strategy $p(z) = z$, in the generic situation when H_0 has distinct moduli, the iterates H_t always converge to an upper triangular matrix H_∞ ; this is because iterates of the unshifted QR algorithm are in one-to-one correspondence with the iterates generated by running simultaneous iteration (a specific form of subspace iteration) for H_0 (or equivalently inverse iteration for H_0^*) on the canonical orthonormal basis; we refer the reader to [56] for details. Combining the unitary similarities accumulated during the iteration, these two facts yield a Schur factorization $A = Q^* H_\infty Q$ of the original matrix, from which the eigenvalues of A can be read off. The unshifted QR iteration does *not* give an efficient algorithm, however, as it is easy to see that convergence can be arbitrarily slow if the ratios of the magnitudes of the eigenvalues of H_0 are close to 1. The role of the shifting strategy is to adaptively improve these ratios and thereby accelerate convergence. The challenge is that this must be done efficiently without prior knowledge of the eigenvalues.

We quantify the speed of convergence of a sequence of iterates of (1) in terms of its δ -decoupling time $\text{dec}_\delta(H_0)$, which is defined as the smallest t at which some subdiagonal entry of H_t satisfies

$$|H_t(i+1, i)| \leq \delta \|H_t\|,$$

where henceforth $\|\cdot\|$ will be used to denote the operator norm. In this context, “rapid” convergence means that $\text{dec}_\delta(H_0)$ is a very slowly growing function of n and $1/\delta$, ideally logarithmic or polylogarithmic. We will refer to a Hessenberg matrix with some $H(i+1, i) = 0$ as *decoupled*.

Remark 1.1 (*Arithmetic Complexity from Decoupling Time*) The motivation for the particular measure of convergence above is that there is a procedure called *deflation* which zeroes out the smallest subdiagonal entry of a δ -decoupled Hessenberg matrix and obtains a nearby block upper triangular matrix, which allows one to pass to subproblems of smaller size incurring a backward error of $\delta \|H_0\|$. Repeating this procedure n times (and exploiting the special structure of Hessenberg matrices to compute the Q_t efficiently) yields an algorithm for computing a triangular T and unitary Q such that $\|H_0 - Q^* T Q\| \leq n\delta \|H_0\|$ in a total of $O(n^3 \text{dec}_\delta(H_0))$ arithmetic operations [58]. Thus, the interesting regime is to take $\delta \ll 1/n$. We provide a full analysis

of the deflation step as well as the total complexity of the algorithm in the subsequent paper [6], but focus solely on decoupling in this paper.

In a celebrated work, Wilkinson [60] proved global convergence⁵ of shifted QR on all *real symmetric* tridiagonal⁶ matrices using the shifting strategy that now carries his name. The linear convergence bound $\text{dec}_\delta(H_0) \leq O(\log(1/\delta))$ for this shifting strategy was then obtained by Dekker and Traub [22] (in the more general setting of Hermitian matrices), and reproven by Hoffman and Parlett [33] using different arguments. Other than these results for Hermitian matrices, there is no known bound on the worst-case decoupling time of shifted QR for any large class of matrices or any other shifting strategy. For nonnormal matrices, it is not even known if there is a shifting strategy which yields global convergence regardless of an effective bound on the decoupling time.⁷ Shifted QR is nonetheless the most commonly used algorithm in practice for the nonsymmetric eigenproblem on dense matrices. The strategies implemented in standard software libraries heuristically converge very rapidly on “typical” inputs, but occasionally examples of nonconvergence are found [21, 38] and dealt with in ad hoc ways.

Accordingly, the main theoretical question concerning shifted QR, which has remained open since the 1960s, is:

Question 1 *Is there a shifting strategy for which the Hessenberg shifted QR iteration provably and rapidly decouples on nonsymmetric matrices?*

Question 1 was asked in various forms e.g. by Parlett [43, 44], Moler [38, 39], Demmel [23, Ch. 4], Higham [32, IV.10], and Smale [50] (who referred to it as a “great challenge”).

The main result of this article (Theorem 1.6) is a positive answer to Question 1 which is quantified in terms of how nonnormal the input matrix is. To be precise, let $\mathbb{H}_B^{n \times n}$ be the set of diagonalizable complex Hessenberg matrices H_0 with eigenvector condition number $\kappa_V(H_0) \leq B$. We exhibit a two parameter family of deterministic shifting strategies $\text{Sh}_{k,B}$ indexed by a degree parameter $k = 2, 4, 8, \dots$ and condition parameter $B \geq 1$ and prove that:

- (i) The strategy $\text{Sh}_{k,B}$ satisfies $\text{dec}_\delta(H_0) \leq O(\log(1/\delta))$ for every $H_0 \in \mathbb{H}_B^{n \times n}$ and $\delta > 0$.
- (ii) $\text{Sh}_{k,B}$ has degree k and can be computed in roughly $O((\log k + B^{\frac{\log k}{k}})kn^2)$ arithmetic operations, which is simply $O(n^2k \log k)$ for the judicious setting $k = \Omega(\log B \log \log B)$.

Thus, the computational cost of the shifting strategy required for convergence increases as the eigenvectors of the input matrix become more and more ill-conditioned,⁸ but the dependence on the eigenvector condition number is very mild.

⁵ i.e., from any initial condition H_0 .

⁶ i.e., arising as the Hessenberg form of symmetric matrices.

⁷ A thorough discussion of related work appears in Sect. 1.2.

⁸ As in many other settings, the condition number of the problem does not only have an effect on the numerical stability, but also on the complexity of the algorithm. We refer the reader to Section 1.1.1 of [5] for a discussion on the condition number of the eigenvalue problem and its relation to $\kappa_V(H)$.

We remark that such a result was not previously known even in the case $B = 1$, which corresponds to normal matrices. Further, as we explain in Remark 1.8, a tiny random perturbation of any $H_0 \in \mathbb{H}^{n \times n}$ is likely to be an element of $\mathbb{H}_B^{n \times n}$ for small B (not depending on H_0). Thus, while our theorem does not give a single shifting strategy which works for all matrices, it does give a strategy which works for a tiny random perturbation of every matrix (with high probability, where “tiny” and “small” must be quantified appropriately). Consequently, the present work can be interpreted in the following three ways:

- (A) *An algorithm with guarantees on a restricted set of inputs.* After fixing the parameter $B \geq 1$, our results provide a shifting strategy which may be run on any input but is only guaranteed to succeed on those inputs H_0 satisfying $\kappa_V(H_0) \leq B$. This may be of use when a priori information about the matrix is given.
- (B) *An algorithm with guarantees for all inputs.* One can adopt a smoothed analysis perspective, in the sense of [51], and further endow the algorithm with a preprocessing step consisting of randomly perturbing the input. This yields a guarantee (with high probability) of success for all inputs and has the notable feature that the running time of the algorithm depends very mildly on the (inverse of the) size of the perturbation. However, in cases when the input matrix possesses some nice structure (e.g. sparsity), this approach has a clear drawback, which points towards investigating the regularization effect of structure-preserving random perturbations.
- (C) *A heuristic for why success occurs “most of the time”.* One can interpret our results as a quantitative statement about the shifting strategy succeeding for “almost all” matrices. Then heuristically (and partially exiting the theoretical outlook of our work) the rounding errors coming from floating point computations could potentially have a regularization effect on the input, ultimately implying convergence even on those inputs that did not satisfy the bound $\kappa_V(H_0) \leq B$.

Motivation and Scope (Theory vs Practice). Before proceeding to formally state the main results of this paper we would like to comment on the motivation and scope of our work. In this regard, it is important to distinguish theory from practice, and clarify that the present paper does not seek to be an immediate prescription for practitioners: more experimentation and theoretical development would be required to perfect the existing (extremely sophisticated) libraries for diagonalization, which throughout the years have been fine-tuned to enhance their performance and efficiency, and constitute an engineering feat.

On the other hand, we do believe that our work greatly advances the theory on the Hessenberg QR algorithm, whose convergence properties (as mentioned above) were mysterious even in the normal case. In particular, this paper provides (in exact arithmetic) a rigorous and conceptual understanding of: (i) the interaction between Ritz values based shifts and exceptional shifts (which was previously understood only in the case of unitary matrices [53–55], and has been exploited in practice in ad hoc ways), and (ii) the impact that ill-conditioned eigenvectors have on the speed of convergence of the algorithm, together with the relevance of higher degree shifts (which are also used in practice) to nonnormality. The subsequent paper [6] further analyzes the numerical instabilities that arise near decoupling when shifted QR is

implemented in finite precision arithmetic, as well as a full analysis of the deflation step, a topic that has been only partially addressed in the past [47, 57] even in the Hermitian setting.

Finally, we also believe that our work has the potential to eventually have a practical impact. As mentioned above, current implementations of the Hessenberg QR algorithm are a result of decades of improvement, where failure cases have been addressed by adding ad hoc shifts, ultimately leading to a complex algorithm which can still potentially fail to converge. For example, [18, pg. 3] states “As implemented in LAPACK 3.1, the new QR algorithm is substantially more complicated...[and] has over 1,300 executable lines of code spread among seven subroutines.” In contrast, the results here provide a conceptually simple shifting strategy which is infallible in the sense of (A) and (B), but which is unlikely to have comparable efficiency to the shifting strategies used in practice. This leads one to wonder if there is a shifting strategy that gets the best of both worlds.

1.1 Statement of Results and Key Ideas

This paper contains two theorems. The first is a warmup to the main theorem, corresponding to the special case $k = 2$, $B = 1$, and states that a certain simple shifting strategy is always rapidly convergent on normal matrices in exact arithmetic, where we assume we can compute square roots exactly.

Theorem 1.2 (Normal Matrices) *There is a degree $k = 2$ deterministic shifting strategy $\text{Sh}_{2,1}$ which ensures that*

$$\text{dec}_\delta(H_0) \leq 4 \log(1/\delta). \quad (2)$$

for all normal Hessenberg H_0 and all $\delta > 0$. The worst case arithmetic complexity of each iteration of $\text{Sh}_{2,1}$ is at most

$$792 \cdot T_{\text{iqr}}(2, n) + O(1)$$

arithmetic operations, where $T_{\text{iqr}}(2, n) \leq 14n^2$ is the arithmetic complexity of an implicit QR step (see Sect. 2).

Key Ideas The proof of Theorem 1.2 appears in Sect. 3. The main challenge in proving such a theorem is that all known shifting strategies based on Ritz values have attractive fixed points which are not decoupled (e.g., certain unitary matrices), and the shifted QR iteration can stagnate (converge slowly) in the neighborhood of such matrices and possibly others. Our main conceptual idea is to carefully design a shifting strategy in such a way that *all of its stagnant states are well-characterized*.⁹ Given our characterization, the shifting strategy then uses another mechanism (certain “exceptional

⁹ The earliest precursor to such a result is Parlett [42], who showed that all of the fixed points (which can be thought of as perfectly stagnant states) of a certain degree 2 shifting strategy are translations of unitary matrices.

shifts”, see Sect. 1.2 for a discussion) to guarantee rapid convergence whenever stagnation occurs. At a technical level, our key innovation is to analyze convergence in terms of certain spectral measures associated with the iterates H_t , which enables our characterization. Our analytic approach differs significantly from previous essentially algebraic (e.g. [22, 33, 53, 60]) and geometric (e.g. [9, 13, 14, 37]) approaches to analyzing the shifted QR algorithm.

Remark 1.3 (*Optimizing the constants in Theorem 1.2*) Our goal in presenting Theorem 1.2 is to elucidate one of the key mechanisms in our analysis by analyzing the simplest instantiation of our shifting strategy, prioritizing clarity of exposition over optimality. Later, in Lemma 4.7 (which applies in a more general setting) we demonstrate that in the normal case the related strategy $\text{Sh}_{4,1}$ achieves 2 using at most 54 calls to a degree 4 implicit QR step in each iteration, instead of the 792 above. Furthermore, as discussed in Remark 3.5, a more meticulous analysis in the normal case can reduce the constant even further, significantly decreasing it from 54 to a smaller number.

Remark 1.4 (*Maintaining Similarities*) All arithmetic complexity bounds in this paper refer only to task of computing the Hessenberg iterates H_t , without keeping track of the accumulated unitary similarities $Q_0 Q_1 \dots Q_t$ between H_t and H_0 . The additional task of computing the similarities is well-understood (e.g. see [59]), and if desired can be achieved with a small increase in the total running time, without changing the asymptotic complexity. We focus on the convergence of the Hessenberg iterates H_t and do not discuss maintaining the similarities further in this paper.

Our second (and main) theorem generalizes Theorem 1.2 to nonnormal matrices. We need the following notion to precisely state it.

θ -Optimal Ritz Values. Like most previously studied shifting strategies, we define $\text{Sh}_{k,B}(H_t)$ in terms of *Ritz values* of the current iterate H_t . In this paper, when working with a Hessenberg matrix, the term Ritz values will be exclusively used to refer to the eigenvalues of its bottom right $k \times k$ corner $(H)_{(k)}$ ¹⁰ they are also characterized variationally as the roots of the degree k monic polynomial p minimizing $\|e_n^* p(H)\|$, where e_n is the last standard basis vector (see Lemma 2.3 for details). Since computing eigenvalues of arbitrary matrices exactly is impossible, we assume access to a method for computing approximate Ritz values, in the sense encapsulated in the following definition.

Definition 1.5 (*θ -Optimal Ritz values and Ritz value finders*) Let $\theta \geq 1$. We call $\mathcal{R} = \{r_1, \dots, r_k\} \subset \mathbb{C}$ a set of *θ -optimal Ritz values* of a Hessenberg matrix H if

$$\left\| e_n^* \prod_{i \leq k} (H - r_i) \right\|^{1/k} \leq \theta \min_{p \in \mathcal{P}_k} \|e_n^* p(H)\|^{1/k}. \tag{3}$$

A *Ritz value finder* is an algorithm $\text{OptRitz}(H, k, \theta)$ that takes as inputs a Hessenberg matrix $H \in \mathbb{C}^{n \times n}$, a positive integer k and an accuracy parameter $\theta > 1$, and outputs a

¹⁰ In general, for a matrix A and a tall matrix with orthonormal columns $Q \in \mathbb{C}^{n \times k}$, the eigenvalues of $Q^* A Q$ are referred to as Ritz values. Here we are only interested in the case when H is Hessenberg and $Q^* = [I_k, 0_{n-k}]$.

set $\mathcal{R} = \{r_1, \dots, r_k\}$ of θ -optimal Ritz values of H whenever the right hand side of (3) is nonzero. Let $T_{\text{OptRitz}}(k, \theta, \delta)$ be the maximum number of arithmetic operations used by $\text{OptRitz}(H, k, \theta)$ over all inputs H such that the right hand side of (3) satisfies¹¹

$$\min_{p \in \mathcal{P}_k} \|e_n^* p(H)\|^{1/k} \geq \delta \|H\|.$$

A Ritz value finder satisfying Definition 1.5 can be efficiently instantiated using polynomial root finders (e.g. [41]) or other provable eigenvalue computation algorithms (e.g. [5, 7]) with guarantees of type $T_{\text{OptRitz}}(\theta, k, \delta) = O(k^c \log(\frac{1}{\delta(\theta-1)}))$. We defer a detailed discussion of numerical issues surrounding this implementation to our companion papers [6, 7]. The subtlety of not being able to compute Ritz values exactly is secondary to the dynamical phenomena which are the focus of this paper, so on first reading of the proofs it is recommended to assume $\theta = 1$ (i.e., Ritz values are computed exactly), even though this is unrealistic when $k > 1$. The theorem below is stated with $\theta = 2$, which is also the parameter setting used in [6].

We now present our main theorem. All logarithms are base 2.

Theorem 1.6 (Nonnormal Matrices) *There is a family of deterministic shifting strategies $\text{Sh}_{k,B}$ (described in Sect. 4) parameterized by degree $k = 2, 4, 8, \dots$ and nonnormality bound $B \geq 1$ with the following properties.*

1. (Rapid Decoupling) *If $H_0 \in \mathbb{H}_B^{n \times n}$, then for every $\delta > 0$, the QR iteration with strategy $\text{Sh}_{k,B}$ satisfies*

$$\text{dec}_\delta(H_0) \leq 4 \log(1/\delta). \tag{4}$$

2. (Cost Per Iteration Before Decoupling) *Given a Ritz value finder $\text{OptRitz}(H, k, \theta)$ with complexity $T_{\text{OptRitz}}(k, \theta, \delta)$, an accuracy parameter $\delta > 0$, and a Hessenberg matrix $H_t \in \mathbb{H}_B^{n \times n}$, computing H_{t+1} given H_t has a cost per iteration of at most*

$$\left(\log k + N_{\text{net}} \left(0.002 B^{-\frac{8 \log k + 4}{k-1}} \right) \right) \cdot T_{\text{qr}}(k, n) + T_{\text{OptRitz}}(k, 2, \delta) + \log k \tag{5}$$

arithmetic operations for all iterations before (4) is satisfied, where $N_{\text{net}}(\epsilon) \leq 4/\epsilon^2$ is number of points in an efficiently computable ϵ -net of the unit disk and $T_{\text{qr}}(k, n) \leq 7kn^2$ is an upper bound on the arithmetic cost of a degree k implicit QR step (see Sect. 2).

The term involving N_{net} captures the the cost of performing certain “exceptional shifts” (see Sect. 1.2) used in the strategy. The tradeoff between the nonnormality of the input matrix and the efficiency of the shifting strategy appears in the cost of the exceptional shift, where it is seen that setting

$$k = \Omega(\log B \log \log B) \tag{6}$$

¹¹ Such a lower bound is needed, since otherwise we could use OptRitz to compute the eigenvalues of $H_{(k)}$ to arbitrary accuracy in finite time.

yields $B^{-\frac{8 \log k + 4}{k-1}} = \Omega(1)$ and a consequent total running time of $O(n^2 k \log k)$ operations per iteration. Note that the bound $B \geq \kappa_V(H_0)$ must be known in advance in order to determine how large a k is needed to make the cost of the exceptional shift small. One may also take k to be a constant independent of B , but this causes the arithmetic complexity of each iteration to depend polynomially on B rather than logarithmically. Note that for normal matrices one may take $k = 2$ and $B = 1$.

Remark 1.7 (Higher Degree Shifts) In exact arithmetic, a QR step with a degree k shift $p(z) = (z - r_1) \dots (z - r_k)$ is identical to a sequence of k steps with degree 1 shifts $(z - r_1), (z - r_2), \dots, (z - r_k)$ (see e.g. [58] for a proof), so any degree k strategy can be simulated by a degree 1 strategy while increasing the iteration count by a factor of k .¹² We choose to present our strategy as higher degree for conceptual clarity. The overall performance of shifting strategies of degrees as high as $k = 180$ has been tested in the past [15, Section 3] and $k = 50$ is often used in practice [35].

Key Ideas. The proof of Theorem 1.6 appears in Sect. 4. The main new difficulty in the nonnormal case is that the iterates H_t can behave chaotically on short time scales,¹³ lacking any kind of obvious algebraic or geometric monotonicity properties (which are present in the normal case, see e.g. [10, 44]). This lack of monotone quantities makes it hard to reason about convergence, as noted by Parlett [44]. For example, consider the family of $n \times n$ matrices:

$$M = \begin{pmatrix} & & & & \beta_n \\ & & & & \\ & \beta_1 & & & \\ & & \beta_2 & & \\ & & & \ddots & \\ & & & & \beta_{n-1} \end{pmatrix} \tag{7}$$

where $\beta_1, \dots, \beta_n \in (0, 1)$. Observe that, for $k \leq n - 1$, the characteristic polynomial of the bottom right corner $M_{(k)}$ is just z^k , so any naïve shifting strategy based on Ritz values will yield the trivial shift. One can verify that a QR step with the trivial strategy applied to M cyclically permutes the β_i , while leaving the zero pattern of M intact. This means that for adversarially chosen β_1, \dots, β_n , the bottom few subdiagonal entries of M —the traditional place to look for monotonicity in order to prove convergence (see e.g. [33])—exhibit arbitrary behavior over a small number of QR steps. At very long time scales of n steps, the behavior becomes periodic and predictable, but there is still no convergence.

We surmount the above difficulty by using *higher degree shifts*. The key insight is that if there is an upperbound on $\kappa_V(H_t)$, then the behavior of a degree $k = \log \kappa_V(H_t)$ QR step is quite predictable, and can be analyzed similarly to the normal case—essentially, this corresponds to k degree 1 steps which is enough to “damp” the transient behavior due to nonnormality (this is articulated precisely in Sect. 4). To see this

¹² This also has some important advantages with regards to numerical stability, which are discussed in [6].

¹³ We measure time not as the number of QR steps, but as the number of QR steps of degree 1, so for example a QR step with a degree k shift corresponds to k time steps.

phenomenon in action, if we impose a bound on $\kappa_V(M)$ in Example (7), it can be seen that the ratios of the β_i cannot be arbitrary and the geometric mean of the bottom $k = \log \kappa_V(M)$ subdiagonal entries of M must remain almost-constant on time intervals of k unshifted QR steps.

At a technical level, the main obstacle is that unlike normal matrices, nonnormal matrices do not have spectral measures or a corresponding continuous functional calculus (which was used crucially in our characterization of stagnation in the proof of Theorem 1.2). A key ingredient in the proof of Theorem 1.6 is a notion of “approximate functional calculus” which allows us to recover enough analytic structure in the nonnormal case to execute the same proof strategy as in the normal case.

Remark 1.8 (*Regularization of κ_V by Random Perturbation*) The *pseudospectral regularization* guarantees from [8] (verifying the conjecture of [20]) imply that every matrix A is $\delta\|A\|$ -close in operator norm to matrices with $\kappa_V = O(n^2/\delta)$. Such a nearby, well-conditioned matrix can be produced (with high probability) by perturbing each entry of A with an independent complex Gaussian of variance δ .¹⁴ After perturbing we can thus (with high probability) take $B = O(n^2/\delta)$ in Theorem 1.6 and set $k = O(\log(n/\delta) \log \log(n/\delta))$, incurring a backward error of δ and yielding a per iteration arithmetic cost of $O(n^2 \log(n/\delta) \log \log^2(n/\delta))$ before δ -decoupling. This approach embraces the fact that the shifted QR algorithm can only in the first place guarantee backward accuracy of the eigenvalues it computes, so there is no harm in using an initial small random perturbation as a “preconditioning” step.

Remark 1.9 (*Numerical Stability, Deflation, and Bit Complexity*) The shifting strategy in Theorem 1.6 can be implemented in floating point arithmetic using $O(k \log(n/\delta))$ bits of precision for the implicit QR steps¹⁵ and $O(k^2 \log^2(n/\delta))$ bits of precision for the Ritz value finder,¹⁶ while preserving both correctness and rapid convergence, with the caveat that the numerical implementation requires using randomization in order to be efficient. This is proved in the companion papers [6, 7], along with a detailed analysis of deflation, yielding a complete algorithm for computing the eigenvalues of a matrix with good bit complexity estimates.

When succeeding, the QR iteration computes the Schur factorization of the input matrix A , which consists of a triangular matrix (whose diagonal entries are the eigenvalues of A), and a unitary matrix (from where the eigenvectors of A can be recovered). The computation of the unitary part is optional, and requires keeping track of the unitary conjugations used throughout the iteration, a task that incurs a higher cost in the running time (e.g. see [59]). In this paper (and the subsequent works [6, 7]) we will focus on the computation of the triangular part (which is the part of the algorithm that was not yet understood), and the running times appearing in the results will refer to exclusively this task. We note however, that a running time for obtaining the full Schur form (i.e. computing the unitary part too), can be obtained directly from our results and will have the same asymptotic running time, just with larger constants.

¹⁴ See also [4, 34] for more general perturbations, including real perturbations, which have a similar effect and [2] for similar results in the complex Gaussian case.

¹⁵ Hence, when a random perturbation is used as a preconditioner, in view of Remark 1.8 the number of bits of precision required becomes $O(\log^2(n/\delta) \log \log(n/\delta))$.

¹⁶ The Ritz value finder is invoked only on small $k \times k$ matrices and incurs a subdominant cost.

1.2 History and Related Work

The literature on shifted QR is vast, so we mention only the most relevant works—in particular, we omit the large body of experimental work and do not discuss the many works on local convergence of shifted QR (i.e., starting from an H_0 which is already very close to decoupling). The reader is directed to the excellent surveys [11, 19, 50] or [29, 45, 59] for a dynamical or numerical viewpoint, respectively, or to the books [23, 30, 52, 58] for a comprehensive treatment. A detailed historical summary appears in [29].

Most of the shifting strategies studied in the literature are a combination of the following three types. The motivation for considering shifts depending on $H_{(k)}$ is closely related to Krylov subspace methods, see e.g. [58]. Below H denotes the current Hessenberg iterate.

1. *k-Francis Shift*. Take $p(z) = \det(z - H_{(k)})$ for some k . The case $k = 1$ is called Rayleigh shift.
2. *Wilkinson Shift*. Take $p(z) = (z - a)$ where a is the root of $\det(z - H_{(2)})$ closer to $H_{(1)}$.
3. *Exceptional Shift*. Let $p(z) = (z - x)$ for some x chosen randomly or arbitrarily, perhaps with a specified magnitude (e.g. $|x| = 1$ for unitary matrices in [25, 53–55]).

Shifting strategies which combine more than one of these through some kind of case analysis are called “mixed” strategies.

Symmetric Matrices. Jiang [26] showed that the geometric mean of the bottom k subdiagonal entries is monotone for the k -Francis strategy in the case of symmetric tridiagonal matrices. Aishima et al. [1] showed that this monotonicity continues to hold for a “Wilkinson-like” shift which chooses $k - 1$ out of k Ritz values. Both of these results yield global convergence on symmetric tridiagonal matrices (without an effective bound on the number of iterations).

Rayleigh Quotient Iteration and Normal Matrices. The behavior of shifted QR is well known to be related to shifted inverse iteration (see e.g. [52]). In particular, the Rayleigh shifting strategy corresponds to a vector iteration process known as Rayleigh Quotient Iteration (RQI). Parlett [44] (building on [17, 40, 46]) showed that RQI converges globally (but without an effective bound) on almost every normal matrix and investigated how to generalize this to the nonnormal case.

Batterson [9] studied the convergence of 2-Francis shifted QR on 3×3 normal matrices with a certain exceptional shift and showed that it always converges. The subsequent work [10] showed that 2-Francis shifted QR converges globally on almost every real $n \times n$ normal matrix (without an effective bound). In Theorem 6 of that paper, it was shown that the same potential that we consider is monotone-decreasing when the k -Francis shift is run on normal matrices, which was an inspiration for our proof of almost-monotonicity for nonnormal matrices.

Nonnormal Matrices. Parlett [42] showed that an unshifted QR step applied to a singular matrix leads to immediate 0-decoupling, taking care of the singularity issue that was glossed over in the introduction, and further proved that all of the fixed points

of an extension of the 2-Francis shifted QR step (for general matrices) are multiples of unitary matrices.

In a sequence of works, Batterson and coauthors investigated the behavior of RQI and 2-Francis on nonnormal matrices from a dynamical systems perspective. Batterson and Smillie [13, 14] showed that there are real matrices such that RQI fails to converge for an open set of real starting vectors. The latter paper also established that RQI exhibits chaotic behavior on some instances, in the sense of having periodic points of infinitely many periods. Batterson and Day [12] showed that 2-Francis shifted QR converges globally and linearly on a certain conjugacy class of 4×4 Hessenberg matrices.

In the realm of periodicity and symmetry breaking, Day [21], building on an example of Demmel, showed that there is an open set of 4×4 matrices on which certain mixed shifting strategies used in the EISPACK library fail to converge rapidly in exact arithmetic; such an example was independently discovered by Moler [38] who described its behavior in finite precision arithmetic. These examples are almost normal in the sense that they satisfy $\kappa_V \leq 2$, so the reason for nonconvergence is symmetry, and our strategy $\text{Sh}_{k,B}$ with modest parameters $k = B = 2$ is guaranteed to converge rapidly on them (in exact arithmetic).

Using topological considerations, Leite et al. [37] proved that no continuous shifting strategy can decouple on every symmetric matrix. Accordingly (in retrospect), the most successful shifting strategy for symmetric matrices, the Wilkinson Shift, is discontinuous in the entries of the matrix and explicitly breaks symmetry when it occurs. Our strategy $\text{Sh}_{k,B}$ is also discontinuous in the entries of the matrix.

Mixed and Exceptional Shifts. Eberlein and Huang [25] showed global convergence (without an effective bound) of a certain mixed strategy for unitary Hessenberg matrices; more recently, the works [53–55] exhibited mixed strategies which converge globally and linearly for unitary Hessenberg matrices with a bound on the rate,¹⁷ but this bound depends on the matrix in a complicated way and is not clearly bounded away from 1. Our strategy $\text{Sh}_{k,B}$ is also a mixed strategy which in a sense combines all three types above. Our choice of exceptional shift was in particular inspired by the work of [25, 54]—the difference is that the size of the exceptional shift is naturally of order 1 in the unitary case, but in the general case it must be chosen carefully at the correct spectral scale.

Higher Degree Shifts. The idea of using higher degree shifts was already present in [22, 27], but was popularized in by Bai and Demmel in [3], who observed that higher order shifts can sometimes be implemented more efficiently than a sequence of lower order ones; see [3, Section 3] for a discussion of various higher order shifting strategies which were considered in the 1980s. More modern approaches use multishifts [15] in combination with other techniques such as aggressive early deflation [16].

The use of higher degree shifts in this paper is purely in order to tame the effects of nonnormality, and *not* to improve efficiency as in the aforementioned previous works. As explained in Remark 1.7, our algorithms can be equivalently described using sequences of single shifts.

¹⁷ We refer to the constant multiplying the error in each iteration of linear convergence as its *rate*.

We defer a detailed discussion of the extensive related work on numerical issues related to shifted QR as well as a comparison to other algorithms for computing eigenvalues (in particular, [2, 5]) to our companion paper [6].

2 Preliminaries and Notation

As mentioned above, the eigenvector condition number of a diagonalizable matrix M is defined as

$$\kappa_V(M) := \inf_{V: M=VDV^{-1}} \|V\| \|V^{-1}\|.$$

We note that, since $\|V\| \|V^{-1}\|$ is invariant under multiplying V by a scalar, the infimum is not changed when restricting to the compact set of eigenvector matrices with $\|V\| \leq 1$, and therefore the infimum is always attained by some matrix V_0 . Moreover, by van der Sluis theorem (see [31, Theorem 7.5]), the condition number of V_0 differs at most by a factor of \sqrt{n} from the eigenvector matrix with unit columns.

Throughout the remainder of the paper, $H = (h_{i,j})_{i,j \in [n]}$ will denote an $n \times n$ upper Hessenberg matrix. As in the introduction, we use

$$H_{(k)} \quad \text{and} \quad \chi_k(z)$$

to denote the lower-right $k \times k$ corner of H and its characteristic polynomial respectively. Following the convention in operator theory, we will write scalar multiples of the identity zI with $z \in \mathbb{C}$ as z , e.g., we will write $z - H$ instead of $zI - H$. As mentioned in the introduction, \mathcal{P}_k will always denote the set of monic polynomials of degree k . All matrix norms are operator norms, denoted by $\|\cdot\|$.

Probability. We use standard probabilistic notation: $\mathbb{P}[\cdot]$ denotes the probability of an event and $\mathbb{E}[\cdot]$ denotes the expectation of a random variable. All probabilities and expectations in this paper are with respect to the random variable “ Z_H ”, defined by (13) in Sect. 3 and more generally by (18) in Sect. 4; the random variable in any probabilistic statement will therefore always be unambiguous from the context. The only probabilistic facts we use are linearity of expectation, Jensen’s inequality, and the Paley-Zygmund inequality.

Implicit QR. We assume black box access to a routine (the implicit QR algorithm) for efficiently performing a QR step in $O(kn^2)$ arithmetic operations (rather than the $O(kn^3)$ operations required by a naive method). Since numerical stability issues are not discussed in this paper, this is the only property that we will use of the implicit QR algorithm.

Definition 2.1 (*Implicit QR Algorithm*) For $k \leq n$, an exact implicit QR algorithm $\text{iqr}(H, p(z))$ takes as inputs an irreducible¹⁸ Hessenberg matrix $H \in \mathbb{C}^{n \times n}$ and a

¹⁸ A Hessenberg matrix H is said to be irreducible if all of its subdiagonal entries are non-zero.

polynomial $p(z) = (z - s_1) \cdots (z - s_k)$ and outputs a Hessenberg matrix \widehat{H} satisfying

$$\widehat{H} = Q^* H Q,$$

where Q is a unitary matrix such that $p(H) = QR$ for some upper triangular matrix R , as well as the number $\|e_n^* p^{-1}(H)\|$ (which appears as the bottom right entry of R) whenever $p(H)$ is invertible. It runs in at most

$$T_{\text{iqr}}(k, n) \leq 7kn^2 \tag{8}$$

operations.

Note that in the above definition we have assumed that the input to `iqr` is an *irreducible* Hessenberg matrix. This will not be a problem throughout the paper in the analysis of the algorithm, since the ultimate goal is to prove an upper bound on the number arithmetic operations needed to achieve δ -decoupling, and by definition *reducible* matrices are δ -decoupled for any $\delta > 0$. We refer the reader to [59, Section 3] for a proof in exact arithmetic of the existence of an efficient implicit QR algorithm.

Potential ψ_k . We will use the geometric mean of the last k subdiagonal entries of the H to track convergence of the Shifted QR iteration, since we are guaranteed δ -decoupling once this quantity is smaller than $\delta\|H\|$. More explicitly, we introduce the following definition.

Definition 2.2 (*Potential function $\psi_k(H)$*) The *potential*¹⁹ $\psi_k(H)$ of H to be

$$\psi_k(H) := |h_{n-k, n-k-1} \cdots h_{n, n-1}|^{\frac{1}{k}}. \tag{9}$$

We record two useful lemmas relating the potential $\psi_k(H)$, the Hessenberg structure of H , and their evolution under the shifted QR iteration. The first gives a variational characterization of the potential (see [52, Theorem 34.1]).

Lemma 2.3 (*Variational Formula for ψ_k*) Let $H \in \mathbb{C}^{n \times n}$ be any Hessenberg matrix. Then, for any k

$$\psi_k(H) = \min_{p \in \mathcal{P}_k} \|e_n^* p(H)\|^{\frac{1}{k}},$$

with the minimum attained for $p = \chi_k$.

Proof Since H is upper Hessenberg, for any polynomial $p \in \mathcal{P}_k$ we have

$$p(H)_{n, n-j} = \begin{cases} p(H_{(k)})_{k, k-j+1} & j = 0, \dots, k-1, \\ h_{n-k, n-k-1} \cdots h_{n, n-1} & j = k, \\ 0 & j \geq k+1. \end{cases}$$

¹⁹ We remark in passing that $\log \psi_k(H)$ is an upper bound on a certain logarithmic potential associated with the Ritz values in the sense of potential theory, though we do not use this relation in this paper (the interested reader can consult [49]). The term ‘‘potential’’ here is alluding to the notion of ‘‘potential function’’ often used in computer science as a tool to track the progress of an algorithm.

Thus for every such p ,

$$\min_{p \in \mathcal{P}_k} \|e_n^* p(H)\| \geq |h_{n-k, n-k-1} \cdots h_{n, n-1}| = \psi_k(H)^k,$$

and the bound will be tight for any polynomial whose application to $H_{(k)}$ zeroes out the last row; by Cayley-Hamilton, the matrix $\chi_k(H_{(k)})$ is identically zero. \square

The second lemma gives a mechanism for proving upper bounds on the potential of $\widehat{H} = \text{iqr}(H, p(z))$ in terms of the shift polynomial p . Consequently, the following quantity will prove useful.

Definition 2.4 ($\tau_p(H)$) For a monic polynomial $p \in \mathcal{P}_k$ define

$$\tau_p(H) := \|e_n^* p(H)^{-1}\|^{-\frac{1}{k}}, \tag{10}$$

when $p(H)$ is invertible, and $\tau_p(H) = 0$ otherwise.

The special case $k = 1$ of the $\tau_p(H)$ quantity has been used to great effect in previous work studying linear shifts (e.g. [33]), and our next lemma shows that it bounds the potential of $\widehat{H} = \text{iqr}(H, p(z))$ for shift polynomials p of arbitrary degree.

Lemma 2.5 (Upper Bounds on $\psi_k(\widehat{H})$) *Let $H \in \mathbb{C}^{n \times n}$ be a Hessenberg matrix, $p(z)$ a monic polynomial of degree k and $\widehat{H} = \text{iqr}(H, p(z))$. Then*

$$\psi_k(\widehat{H}) \leq \tau_p(H).$$

Proof Assume first that $p(H)$ is singular. In this case for any QR decomposition $p(H) = QR$, the entry $R_{n,n} = 0$, and because $p(\widehat{H}) = Q^* p(H) Q = RQ$, the last row of $p(\widehat{H})$ is zero as well. In particular $\psi_k(\widehat{H}) = |p(\widehat{H})_{1, k+1}|^{\frac{1}{k}} = 0 = \tau_p(H)$. When $p(H)$ is invertible, applying Lemma 2.3 and using repeatedly that Q is unitary, R is triangular, and $p(H) = QR$,

$$\begin{aligned} \psi_k(\widehat{H})^k &\leq \|e_n^* p(\widehat{H})\| = \|e_n^* Q^* p(H)\| = \|e_n^* R\| = \|e_n^* R^{-1} Q^*\|^{-1} \\ &= \|e_n^* p(H)^{-1}\|^{-1} = \tau_p(H)^k. \end{aligned}$$

\square

Lemma 2.5 ensures that given H , we can reduce the potential with an implicit QR step by producing a polynomial p with $\tau_p(H) = \|e_n^* p(H)^{-1}\|^{-\frac{1}{k}} \leq (1 - \gamma) \psi_k(H)$. Note that if the roots of $p(z)$ are close to the eigenvalues of H , then we expect $\|e_n^* p(H)^{-1}\|$ to be large, and therefore $\tau_p(H)$ to be small, which articulates that shifts that are close to the eigenvalues accelerate convergence.

3 Normal Matrices

In this short section we describe the simplest shifting strategy in our family, $\text{Sh}_{2,1}$, and prove that it enjoys global linear convergence with a uniform rate²⁰ for all normal matrices, improving the qualitative results of [46]. The strategy is inspired by the Wilkinson shift in that it chooses the shift to be one of the roots of χ_2 the characteristic polynomial of $H_{(2)}$, but it does so in a "greedy" manner which is based on the quantity τ defined in (10). If one of these shifts doesn't make adequate progress towards convergence, the strategy resorts to certain carefully chosen exceptional shifts, one of which is guaranteed to do so. We track convergence of the strategy using the potential $\psi_2(H)$ defined in (9), which is simply the geometric mean of the bottom two subdiagonal entries of H .²¹

We will heavily use the functional calculus for normal matrices (a standard tool from functional analysis [48, Chapter VII]) to describe and analyze our strategy. Recall that if H has eigenvalues $\Lambda(H) = \{\lambda_1, \dots, \lambda_n\} \subset \mathbb{C}$ and corresponding orthonormal eigenvectors $v_1, \dots, v_n \in \mathbb{C}^n$, then for any analytic function f defined in a neighborhood of $\Lambda(H)$:

$$\int |f(z)|^2 d\mu(z) = \|e_n^* f(H)\|_2^2, \tag{11}$$

where

$$\mu := \sum_{i \leq n} |e_n^* v_i|^2 \delta_{\lambda_i} \tag{12}$$

is the spectral (probability) measure corresponding to e_n . We use probabilistic instead of linear algebraic notation to make our proofs more transparent: letting Z_H denote the random variable taking values in $\{\lambda_1, \dots, \lambda_n\}$ with distribution μ as above, (11) can be rewritten as

$$\mathbb{E}[|f(Z_H)|^2] = \|e_n^* f(H)\|_2^2. \tag{13}$$

Remark 3.1 (Probabilistic Notation) We stress that our shifting strategy is deterministic and we use random variables for notational convenience only. The proofs could be equivalently written in terms of inner products involving the eigenvectors of H , but this would obscure the convexity and probabilistic inequalities that drive them.

The shifting strategy is presented below. Recall that an ϵ -net of a compact set $X \subset \mathbb{C}$ is a finite set of points $\mathcal{N} \subset X$ such that

$$\max_{x \in X} \min_{y \in \mathcal{N}} |x - y| \leq \epsilon.$$

²⁰ i.e., the error is multiplied by a fixed constant in each iteration.

²¹ The exact same potential function was used by Hoffman and Parlett [33] to analyze the Wilkinson shift on symmetric matrices.

Note that the equality of the last two expressions in Line 1 relies on (13).

Sh_{2,1}, “Greedy Wilkinson Shift for Normal Matrices”

Input: Hessenberg H , decoupling rate $\gamma \in (0, 1)$
Output: Hessenberg \widehat{H}
Requires: $0 < \psi_2(H)$
Ensures: $\psi_2(\widehat{H}) \leq (1 - \gamma)\psi_2(H)$.

1. Let $p(z) = (z - r_1)(z - r_2) = \det(z - H_{(2)})$. Choose

$$r = \arg \max_{i=1,2} \|e_n^*(H - r_i)^{-2}\| = \arg \max_{i=1,2} \mathbb{E}[|Z_H - r_i|^{-4}]^{1/2},$$
 computing these quantities using $\text{iqr}(H, (z - r_i)^2)$ for $i = 1, 2$.
2. If $\psi_2(\text{iqr}(H, (z - r)^2)) \leq (1 - \gamma)\psi_2(H)$, output $\widehat{H} = \text{iqr}(H, (z - r)^2)$
3. Else, let \mathcal{S} be a $\epsilon\psi_2(H)$ -net of $D(r, \sqrt{3}\psi_2(H))$ with $\epsilon = (1 - \gamma)^2/\sqrt{27}$.
4. For each $s \in \mathcal{S}$, if $\psi_2(\text{iqr}(H, (z - s)^2)) \leq (1 - \gamma)\psi_2(H)$, output $\widehat{H} = \text{iqr}(H, (z - s)^2)$

The key observation is that if Sh_{2,1} fails to significantly decrease ψ_2 in Line 2, then the spectral measure of H with respect to e_n must be significantly supported on a disk of radius roughly $\psi_2(H)$ centered at r .

Lemma 3.2 (Stagnation Implies Support, Normal Case) *Let $\gamma \in (0, 1)$ and let r be the Greedy Wilkinson shift for an upper Hessenberg matrix H . If*

$$\psi_2\left(\text{iqr}(H, (z - r)^2)\right) \geq (1 - \gamma)\psi_2(H) > 0 \tag{14}$$

then for every $t \in (0, 1)$:

$$\mathbb{P}\left[|Z_H - r| \leq \frac{1}{\sqrt{t}}\psi_2(H)\right] \geq (1 - t)^2(1 - \gamma)^4 \tag{15}$$

Proof Observe that $H - r$ is invertible since otherwise, for $\widehat{H} = \text{iqr}(H, (z - r)^2)$, we would have $\psi_2(\widehat{H}) = 0$ by Lemma 2.5. Our assumption implies that:

$$\begin{aligned}
 (1 - \gamma)\psi_2(H) &\leq \psi_2(\widehat{H}) && \text{hypothesis} \\
 &\leq \tau_{(z-r)^2}(H) && \text{Lemma 2.5} \\
 &= \|e_n^*(H - r)^{-2}\|^{-1/2} && \text{definition} \\
 &= \left(\mathbb{E} \left[\frac{1}{|Z_H - r|^4} \right] \right)^{-1/4} && \text{by (13)} \\
 &\leq \left(\mathbb{E} \left[\frac{1}{|Z_H - r|^2} \right] \right)^{-1/2} && \text{Jensen, } x \mapsto x^2 \\
 &\leq \left(\frac{1}{2} \sum_{i=1}^2 \mathbb{E} \left[\frac{1}{|Z_H - r_i|^2} \right] \right)^{-1/2} && \text{choice of } r \text{ in Line 1 and (13)} \\
 &= \left(\mathbb{E} \left[\frac{1}{2} \sum_{i=1}^2 \frac{1}{|Z_H - r_i|^2} \right] \right)^{-1/2} && \text{Fubini} \\
 &\leq \left(\mathbb{E} \left[\frac{1}{|Z_H - r_1||Z_H - r_2|} \right] \right)^{-1/2} && \text{AM/GM} \\
 &= \left(\mathbb{E} \left[\frac{1}{|\chi_2(Z_H)|} \right] \right)^{-1/2} && \text{definition of } \chi_2 \\
 &\leq (\mathbb{E}[|\chi_2(Z_H)|])^{1/2} && \text{Jensen, } x \mapsto 1/x, x > 0 \\
 &\leq \left(\mathbb{E} \left[|\chi_2(Z_H)|^2 \right] \right)^{1/4} && \text{Jensen, } x \mapsto x^2 \\
 &= \psi_2(H). && \text{Lemma 2.3 and (13)}
 \end{aligned}$$

Thus, all quantities appearing in the above chain of all inequalities lie within a multiplicative factor of $(1 - \gamma)$ of each other. Rearranging and examining the fourth and fifth lines, we obtain the following useful bound:

$$\mathbb{E} \left[|Z_H - r|^{-4} \right] \leq \frac{1}{(1 - \gamma)^4} \left(\mathbb{E}[|Z_H - r|^{-2}] \right)^2. \quad (16)$$

Note that the above inequality is in some sense a “reverse Jensen” inequality, since the actual Jensen inequality for the function $x \mapsto x^2$ yields $\mathbb{E}[|Z_H - r|^{-4}] \geq \mathbb{E}[|Z_H - r|^{-2}]^2$. So (16) is articulating that when stagnation happens, the second and fourth moments of $|Z_H - r|^{-1}$ are close to each other, and therefore we should expect concentration for the random variable $|Z_H - r|^{-1}$ (equivalently of $|Z_H - r|$). Explicitly,

we now have

$$\begin{aligned}
 & \mathbb{P} \left[|Z_H - r| \leq \frac{1}{\sqrt{t}} \psi_2(H) \right] \\
 &= \mathbb{P} \left[|Z_H - r|^{-2} \geq \frac{t}{\psi_2(H)^2} \right] \\
 &\geq \mathbb{P} \left[|Z_H - r|^{-2} \geq t \mathbb{E}[|Z_H - r|^{-2}] \right] && \text{since } \tau_{(z-r)^2}(H) \leq \psi_2(H) \\
 &\geq (1-t)^2 \frac{\mathbb{E}[|Z_H - r|^{-2}]^2}{\mathbb{E}[|Z_H - r|^{-4}]} && \text{Paley-Zygmund} \\
 &\geq (1-t)^2 (1-\gamma)^4 && \text{by (16),}
 \end{aligned}$$

establishing (15). □

Using this lemma we now show that $\text{Sh}_{2,1}$ satisfies its guarantees. The following theorem implies Theorem 1.2 by setting $\gamma = 0.2$ and calculating $\frac{12 \cdot 27}{(4/5)^4} \leq 792$.

Theorem 3.3 *Let $\gamma \in (0, 1)$. The strategy $\text{Sh}_{2,1}$ ensures that*

$$\psi_2(\hat{H}) \leq (1 - \gamma)\psi_2(H).$$

The worst case complexity of $\text{Sh}_{2,1}$ is $2 + |\mathcal{S}| \leq 2 + \frac{12 \cdot 27}{(1-\gamma)^4}$ calls to $\text{iqr}(H, (z - (\cdot))^2)$ plus a constant number of arithmetic operations to compute r_1 and r_2 in Line 1.

Proof Suppose Line 2 does not succeed in reducing the potential $\psi_2 = \psi_2(H)$ by a factor of $(1 - \gamma)$. By Lemma 3.2, for $t \in (0, 1)$ to be chosen later:

$$\mathbb{P} \left[|Z_H - r| \leq \psi_2/\sqrt{t} \right] \geq (1-t)^2(1-\gamma)^4.$$

Let \mathcal{S} be an $\epsilon\psi_2$ -net of $D(r, \psi_2/\sqrt{t})$ for ϵ chosen as in Line 3; an elementary packing argument implies that such a net exists with

$$|\mathcal{S}| \leq \frac{\text{area}(D(r, \psi_2/\sqrt{t}))}{\text{area}(D(r, \epsilon\psi_2/2))} = 4/t\epsilon^2. \tag{17}$$

Choose a point $s \in \mathcal{S}$ satisfying

$$\mathbb{P}[Z_H \in D(s, \epsilon\psi_2)] \geq \frac{1}{|\mathcal{S}|}(1-t)^2(1-\gamma)^2,$$

which must exist since

$$D(r, \psi_2/\sqrt{t}) \subset \bigcup_{s \in \mathcal{S}} D(s, \epsilon\psi_2).$$

We then have

$$\begin{aligned} \tau_{(z-s)^2}^{-4}(H) &= \mathbb{E} \left[\frac{1}{|Z_H - s|^4} \right] \geq \mathbb{P}[|Z_H - s| \leq \epsilon\psi_2] \cdot \frac{1}{\epsilon^4\psi_2^4} \\ &\geq \frac{1}{|\mathcal{S}|} (1-t)^2 (1-\gamma)^4 \frac{1}{\epsilon^4\psi_2^4} \\ &\geq \frac{t\epsilon^2}{4} (1-t)^2 (1-\gamma)^4 \frac{1}{\epsilon^4\psi_2^4} \\ &= \frac{(1-\gamma)^4}{27 \cdot \epsilon^2 \psi_2^4}, \end{aligned}$$

choosing $t = 1/3$ to maximize the right hand side in the penultimate line. This ultimately yields

$$\tau_{(z-s)^2}(H) \leq \frac{27^{1/4} \sqrt{\epsilon}}{1-\gamma} \psi_2(H),$$

which by our choice of ϵ in Line 3 and Lemma 2.5 implies:

$$\psi_2(\text{iqr}(H, (z-s)^2)) \leq \tau_{(z-s)^2}(H) \leq (1-\gamma)\psi_2(H),$$

as advertised.

The bound on the complexity follows by simply counting the number of calls to iqr in Lines 2 and 4 and using the estimate $|\mathcal{S}| \leq 4 \cdot 3/\epsilon^2$. \square

Remark 3.4 (*Choice of γ*) The decoupling rate γ may be viewed as a tuning parameter which trades off the worst case complexity of a single step of the strategy against the worst case total number of steps: if γ is larger then each step is guaranteed to make more progress, but the cost of performing exceptional shift is higher as the required ϵ -net is larger.

Remark 3.5 (*Improving the Disk to an Annulus and optimizing Theorem 3.3*) Inequality (15) is a tail bound of the random variable $|Z_H - r|$. We note that the other tail can be controlled too via Markov's inequality and the upper bound on $\mathbb{E}[|Z_H - r|^{-4}]$ obtained in the proof of Lemma 3.2. Then, the control on both tails yields that the distribution of Z_H has significant mass on a thin annulus (the inner and outer radii are almost the same) around r .

It is instructive to note that when $\gamma = 0$, this annulus becomes of width 0, yielding that Z_H is fully supported on a circle with center r and radius $\psi_k(H)$, proving that in this case $\frac{1}{\psi_k(H)}(H - r)$ is in fact a unitary matrix, which recovers Parlett's characterization of fixed points [42] for the Greedy Wilkinson shift.

In any case, controlling both tails allows one to reduce the search performed by the exceptional shifts from a disk to an annulus, significantly decreasing the size of the net $|\mathcal{S}|$ considered in Theorem 3.3. Moreover, one can think of taking an optimal net on the relevant region (instead of just using a packing argument to upper bound the size of the net), using higher degree shifts (as discussed in Remark 1.3), and refining the arguments in Theorem 3.3 to further reduce the complexity of the exceptional shift, which may be relevant in practical implementations. We omit such optimization in this paper for the sake of simplicity.

4 Nonnormal Matrices

In this section we generalize the strategy $\text{Sh}_{2,1}$ for normal matrices to a family $\text{Sh}_{k,B}$ with provably rapid convergence on not necessarily normal matrices satisfying $\kappa_V(H) \leq B$ for some given $B \geq 1$, and prove the main Theorem 1.6 of this paper.

It is instructive to note that the only way in which normality was used in the proof of Theorem 3.3 is the existence of the spectral measure (12) and corresponding functional calculus (13), which crucially implied that

$$\mathbb{E}[|q(Z_H)|^{-2}]^{1/2} = \|e_n^* q(H)^{-1}\|$$

for quadratic polynomials q , guiding our choice of shift and ultimately enabling the proof of Lemma 3.2. This fact is no longer true in the nonnormal case. Our main idea is that when $\kappa_V(H) < \infty$, there is an “approximate functional calculus” which can be used as an effective substitute, provided that we consider shift polynomials q of appropriately large degree. We will heavily use the following construction throughout this section.

Definition 4.1 (*Approximate Functional Calculus*) Assume that $H = VDV^{-1}$ is diagonalizable, with V chosen²² so that $\|V\| = \|V^{-1}\| = \sqrt{\kappa_V(H)}$ and D a diagonal matrix with $D_{i,i} = \lambda_i$, the eigenvalues of H . Define Z_H to be the random variable supported on the eigenvalues of H with distribution

$$\mathbb{P}[Z_H = \lambda_i] = \frac{|e_n^* V e_i|^2}{\|e_n^* V\|^2}. \tag{18}$$

Note that $\mathbb{P}[Z_H = \lambda_i] = 1$ exactly when e_n^* is a left eigenvector with eigenvalue λ_i , and that when H is normal, the distribution of Z_H is the spectral measure of H associated to e_n^* , so definition 4.1 generalizes (13).

²² In the event that there are multiple such choices of V it does not matter which we choose, only that it remains fixed throughout the analysis.

Lemma 4.2 For any upper Hessenberg H and complex function f whose domain includes the eigenvalues of H ,

$$\frac{\|e_n^* f(H)\|}{\kappa_V(H)} \leq \mathbb{E} \left[|f(Z_H)|^2 \right]^{\frac{1}{2}} \leq \kappa_V(H) \|e_n^* f(H)\|.$$

Proof By the definition of Z_H above,

$$\mathbb{E} \left[|f(Z_H)|^2 \right]^{\frac{1}{2}} = \frac{\|e_n^* f(H)V\|}{\|e_n^* V\|} \leq \|e_n^* f(H)\| \|V\| \|V^{-1}\| = \|e_n^* f(H)\| \kappa_V(H),$$

and the left hand inequality is analogous. □

The upshot of this lemma is that if $q \in \mathcal{P}_k$, then $\tau_q(H) = \|e_n^* q(H)^{-1}\|^{1/k}$ approximates $\mathbb{E}[|q(Z_H)|^{-2}]^{1/2k}$ up to a factor of $\kappa_V(H)^{1/k}$, which is close to 1 if we choose $k \gg \log \kappa_V(H)$. Thus, by choosing k large enough we can obtain accurate information about Z_H by examining the observable quantities $\|e_n^* f(H)\|^{1/k}$, which enables a precise understanding of convergence and a generalization of Lemma 3.2 to the nonnormal case. This motivates the use of a higher degree shifting strategy as a way to deal with nonnormality. Since the iterates are all unitarily similar, κ_V is preserved with each iteration, so the k required is an invariant of the algorithm. Thus the use of a sufficiently high-degree shifting strategy is both an essential feature and unavoidable cost of our approach.

In the remainder of this section we describe and analyze the shifting strategy $\text{Sh}_{k,B}$. The proof of the main result appears in Sect. 4.2. The structure of the proof is similar to that of $\text{Sh}_{2,1}$, but with three important differences: (i) We work with polynomials of degree $k \approx \log B \log \log B$ rather than 2 in light of the above discussion. This requires settling for approximate (i.e., θ -optimal) rather than exact Ritz values when $k \geq 5$, even in exact arithmetic. (ii) The “greedy” choice in $\text{Sh}_{2,1}$ cannot be made exactly due to the absence of (13). We introduce the notion of a “promising Ritz value” (Sect. 4.1) as an approximate surrogate for the “greedy” choice with similar properties, and describe an efficient procedure for finding such a Ritz value (Sect. 4.3). (iii) All of the proofs involve carrying around approximation factors arising from the use of Lemma 4.2, θ -optimality, and promising Ritz values. The required exceptional shift (analyzed in Sect. 4.4) is correspondingly larger.

Notation and Constants. $B \geq \kappa_V(H)$ denotes an upper bound on its eigenvector condition number and $k \geq 2$ a power of two, which the reader may consider for concreteness to be on the order of $\log B \log \log B$; all logarithms will be taken base two for simplicity.

Note that the relationship (6) between k and B is *not* required for the proof of potential reduction, but impacts the cost of performing each iteration. The table below collates notation and constants which will appear throughout this section.

Symbol	Meaning	Typical scale
H	Upper Hessenberg matrix	
Z_H	Random variable in Definition 4.1	
B	Eigenvector condition bound	$B \geq \kappa_V(H)$
k	Shift degree	$O(\log B \log \log B)$
δ	Decoupling parameter	
γ	Decoupling rate	0.2
θ	Approximation parameter for Ritz values	2
α	Promising Ritz value parameter	$B^{4k-1} \log k = 1 + o(1)$

4.1 Promising Ritz Values and Almost Monotonicity of the Potential

In the same spirit as Wilkinson’s shift, which chooses a particular Ritz value (out of two), but using a different criterion, our shifting strategy will begin by choosing a Ritz value (out of k) that has the following property for some $\alpha \geq 1$. This is a generalization of the “greedy” Wilkinson shift considered in Sect. 3.

Definition 4.3 (*α -promising Ritz value*) Let $\alpha \geq 1$, $\mathcal{R} = \{r_1, \dots, r_k\}$ be a set of θ -approximate Ritz values for H , and $p(z) = \prod_{i=1}^k (z - r_i)$. We say that $r \in \mathcal{R}$ is *α -promising* if

$$\mathbb{E} \left[\frac{1}{|Z_H - r|^k} \right] \geq \frac{1}{\alpha^k} \mathbb{E} \left[\frac{1}{|p(Z_H)|} \right]. \tag{19}$$

Note that there is at least one 1-promising Ritz value in every set of approximate Ritz values, since

$$\frac{1}{k} \sum_{i=1}^k \mathbb{E} \left[\frac{1}{|Z_H - r_i|^k} \right] = \mathbb{E} \left[\frac{1}{k} \sum_{i=1}^k \frac{1}{|Z_H - r_i|^k} \right] \geq \mathbb{E} \left[\frac{1}{|p(Z_H)|} \right] \tag{20}$$

by linearity of expectation and AM/GM. The notion of α -promising Ritz value is a relaxation which can be computed efficiently from the entries of H (in fact, as we will explain in Sect. 4.3, using a small number of implicit QR steps with Francis-like shifts of degree $k/2$).

As a warm-up for the analysis of the shifting strategy, we will first show that if $k \gg \log \kappa_V(H)$ and r is a promising Ritz value, the potential is *almost monotone* under the shift $(z - r)^k$. This articulates the phenomenon observed in Example (7) and suggests that promising Ritz values should give rise to good polynomial shifts. Monotonicity is not actually used in the proof of our main theorem, which instead relies on the closely related property (21) established below.

Lemma 4.4 (Almost-monotonicity and Moment Comparison) *Let $\mathcal{R} = \{r_1, \dots, r_k\}$ be a set of θ -optimal Ritz values, as in Definition 1.5, and assume that $r \in \mathcal{R}$ is*

α -promising. If $\widehat{H} = \text{iqr}(H, (z - r)^k)$ then

$$\psi_k(\widehat{H}) \leq \kappa_V(H)^{\frac{2}{k}} \alpha \theta \psi_k(H),$$

and moreover

$$\mathbb{E} \left[|Z_H - r|^{-2k} \right] \geq \mathbb{E} \left[|Z_H - r|^{-k} \right]^2 \geq \frac{1}{\kappa_V(H)^2 (\alpha \theta \psi_k(H))^{2k}}. \tag{21}$$

Proof Let $p(z) = \prod_{i=1}^k (z - r_i)$. The claim follows from the following chain of inequalities:

$$\begin{aligned} \sqrt{\mathbb{E} \left[|Z_H - r|^{-2k} \right]} &\geq \mathbb{E} \left[|Z_H - r|^{-k} \right] && \text{Jensen, } x \mapsto x^2 \\ &\geq \frac{1}{\alpha^k} \mathbb{E} [|p(Z_H)|^{-1}] && r \text{ is } \alpha - \text{promising} \\ &\geq \frac{1}{\alpha^k} \frac{1}{\mathbb{E} [|p(Z_H)|]} && \text{Jensen, } x \mapsto 1/x, x > 0 \\ &\geq \frac{1}{\alpha^k} \frac{1}{\sqrt{\mathbb{E} [|p(Z_H)|^2]}} && \text{Jensen, } x \mapsto x^2 \\ &\geq \frac{1}{\alpha^k} \frac{1}{\|e_n^* p(H)\| \kappa_V(H)} && \text{Lemma 4.2} \\ &\geq \frac{1}{\alpha^k} \frac{1}{\theta^k \|e_n^* \chi_k(H)\| \kappa_V(H)} && \text{Definition 1.5 of } \theta - \text{optimal} \\ &= \frac{1}{\alpha^k} \frac{1}{\theta^k \psi_k(H)^k \kappa_V(H)} && \text{Lemma 2.3.} \end{aligned}$$

This already shows (21). For the other claim, rearrange both extremes of the above inequality to get

$$\begin{aligned} \alpha \theta \kappa_V(H)^{\frac{1}{k}} \psi_k(H) &\geq \mathbb{E} \left[|Z_H - r|^{-2k} \right]^{-\frac{1}{2k}} \\ &\geq \frac{\tau_{(z-r)^k}(H)}{\kappa_V(H)^{\frac{1}{k}}} && \text{Lemma 4.2} \\ &\geq \frac{\psi_k(\widehat{H})}{\kappa_V(H)^{\frac{1}{k}}} && \text{Lemma 2.5} \end{aligned}$$

which concludes the proof. □

In Sect. 4.2, we will see that when the shift associated with a promising Ritz value does not reduce the potential, Lemma 4.4 can be used to provide a two-sided bound on the quantities $\mathbb{E}[|Z_H - r|^{-2k}]$ and $\mathbb{E}[|Z_H - r|^{-k}]^2$. This is the main ingredient needed to obtain information about the distribution of Z_H when potential reduction is not achieved.

4.2 The Shifting Strategy

In this section we specify the shifting strategy $\text{Sh}_{k,B}$ and prove Theorem 1.6. An important component of our shifting scheme, presented in detail in Sect. 4.3, is a simple subroutine, “Find,” guaranteed to produce an α -promising Ritz value with $\alpha = \kappa_V(H)^{4k^{-1} \log k}$. Guarantees for this subroutine are stated in the lemma below and proved in Sect. 4.3.

Lemma 4.5 (Guarantees for Find) *The subroutine Find specified in Sect. 4.3 produces a $\kappa_V(H)^{4k^{-1} \log k}$ -promising Ritz value, using at most $12k \log kn^2 + \log k$ arithmetic operations.*

Our strategy is then built around the following dichotomy, which crucially uses the α -promising property: in the event that a degree k implicit QR step with the α -promising Ritz value output by Find does *not* achieve potential reduction, we show that there is a modestly sized set of exceptional shifts, one of which is *guaranteed* to achieve potential reduction. These exceptional shifts are constructed by the procedure “Exc” described in Sect. 4.4. The overall strategy is specified below.

$\text{Sh}_{k,B}$

Input: Hessenberg H and a set \mathcal{R} of θ -approximate Ritz values of H
Output: Hessenberg \widehat{H} , which will be the next matrix in the iteration.
Requires: $0 < \psi_k(H)$ and $\kappa_V(H) \leq B$
Ensures: $\psi_k(\widehat{H}) \leq (1 - \gamma)\psi_k(H)$ and $\kappa_V(\widehat{H}) \leq B$

1. $r \leftarrow \text{Find}(H, \mathcal{R})$
2. If $\psi_k(\text{iqr}(H, (z - r)^k)) \leq (1 - \gamma)\psi_k(H)$, output $\widehat{H} = \text{iqr}(H, (z - r)^k)$
3. Else, $\mathcal{S} \leftarrow \text{Exc}(H, r, B)$
4. For each $s \in \mathcal{S}$, if $\psi_k(\text{iqr}(H, (z - s)^k)) \leq (1 - \gamma)\psi_k(H)$, output $\widehat{H} = \text{iqr}(H, (z - s)^k)$

The failure of line (2) of $\text{Sh}_{k,B}$ to reduce the potential gives useful quantitative information about the distribution of Z_H , articulated in the following lemma. This will then be used to design the set \mathcal{S} of exceptional shifts produced by Exc in line (3) and prove that at least one of them makes progress in line (4).

Lemma 4.6 (Stagnation Implies Support) *Let $\gamma \in (0, 1)$ and $\theta \geq 1$, and let $\mathcal{R} = \{r_1, \dots, r_k\}$ be a set of θ -approximate Ritz values of H . Suppose $r \in \mathcal{R}$ is α -promising and assume*

$$\psi_k\left(\text{iqr}(H, (z - r)^k)\right) \geq (1 - \gamma)\psi_k(H) > 0. \tag{22}$$

Then Z_H is well-supported on an disk of radius approximately $\alpha\psi_k(H)$ centered at r in the following sense: for every $t \in (0, 1)$:

$$\mathbb{P} \left[|Z_H - r| \leq \theta\alpha \left(\frac{\kappa_V(H)}{t} \right)^{\frac{1}{k}} \psi_k(H) \right] \geq (1-t)^2 \frac{(1-\gamma)^{2k}}{\alpha^{2k}\theta^{2k}\kappa_V(H)^4}. \tag{23}$$

Proof Observe that $H - r$ is invertible since otherwise, for $\widehat{H} = \text{iqr}(H, (z-r)^k)$, we would have $\psi_k(\widehat{H}) = 0$ by Lemma 2.5. Our assumption implies that that:

$$\begin{aligned} (1-\gamma)\psi_k(H) &\leq \psi_k(\widehat{H}) && \text{hypothesis} \\ &\leq \tau_{(z-r)^k}(H) && \text{Lemma 2.5} \\ &= \|\varrho_n^*(H-r)^{-k}\|^{-\frac{1}{k}} && \text{definition} \\ &\leq \left(\frac{\kappa_V(H)}{\mathbb{E}[|Z_H - r|^{-2k}]^{\frac{1}{2}}} \right)^{1/k} && \text{Lemma 4.2.} \end{aligned}$$

Rearranging and using (21) from Lemma 4.4 we get

$$\begin{aligned} \frac{\kappa_V(H)^2}{(1-\gamma)^{2k}\psi_k(H)^{2k}} &\geq \mathbb{E}[|Z_H - r|^{-2k}] \geq \mathbb{E}[|Z_H - r|^{-k}]^2 \\ &\geq \frac{1}{\alpha^{2k}\theta^{2k}\psi_k(H)^{2k}\kappa_V(H)^2}, \end{aligned} \tag{24}$$

which upon further rearrangement yields the ‘‘reverse Jensen’’ type bound (note that for the function $x \mapsto x^2$, Jensen’s inequality yields the complementary $\mathbb{E}[|Z_H - r|^{-k}]^2 \leq \mathbb{E}[|Z_H - r|^{-2k}]$):

$$\frac{\mathbb{E}[|Z_H - r|^{-2k}]}{\mathbb{E}[|Z_H - r|^{-k}]^2} \leq \left(\frac{\alpha\theta}{(1-\gamma)} \right)^{2k} \kappa_V(H)^4. \tag{25}$$

We now have

$$\begin{aligned} &\mathbb{P} \left[|Z_H - r| \leq \frac{\alpha}{t^{1/k}} \theta \psi_k(H) \kappa_V^{1/k} \right] \\ &= \mathbb{P} \left[|Z_H - r|^{-k} \geq t \frac{1}{\alpha^k \theta^k \psi_k(H)^k \kappa_V} \right] \\ &\geq \mathbb{P} \left[|Z_H - r|^{-k} \geq t \mathbb{E}[|Z_H - r|^{-k}] \right] && \text{by (24)} \\ &\geq (1-t)^2 \frac{\mathbb{E}[|Z_H - r|^{-k}]^2}{\mathbb{E}[|Z_H - r|^{-2k}]} && \text{Paley-Zygmund} \\ &\geq (1-t)^2 \frac{(1-\gamma)^{2k}}{\alpha^{2k}\theta^{2k}\kappa_V(H)^4} && \text{by (25),} \end{aligned}$$

establishing (23), as desired. □

In Sect. 4.4, we will use Lemma 4.6 to prove the following guarantee on Exc.

Lemma 4.7 (Guarantees for Exc) *The subroutine Exc specified in Sect. 4.4 produces a set \mathcal{S} of exceptional shifts, one of which achieves potential reduction. If $\theta \leq 2$, $\gamma = 0.2$, and $\alpha = B^{4 \log k/k}$, then both the number of degree k iqr calls required for Exc, and the size of \mathcal{S} , are at most*

$$N_{\text{net}} \left(0.002 B^{-\frac{8 \log k + 4}{k}} \right),$$

where $N_{\text{net}}(\epsilon) = O(\epsilon^{-2})$ denotes number of points in an efficiently computable ϵ -net of the unit disk. In the normal case, taking $B = \alpha = \theta = 1$, $k = 4$, $\gamma = 0.2$, the arithmetic operations required and the size of $|\mathcal{S}|$ are both bounded by 50.

We are now ready to prove Theorem 1.6.

Proof of Theorem 1.6 *Rapid convergence.* In the event that we choose a α -promising Ritz value in step (1) that does not achieve potential reduction in step (2), Lemma 4.7 then guarantees we achieve potential reduction in (3). Thus each iteration decreases the potential by a factor of at least $(1 - \gamma)$, and since $\psi_k(H_0) \leq \|H\|$ we need at most

$$\frac{\log(1/\delta)}{\log(1/(1 - \gamma))} \leq 4 \log(1/\delta)$$

iterations before $\psi_k(H_i) \leq \delta \|H_0\|$, which in particular implies δ -decoupling.

Arithmetic Complexity. Computing a full set \mathcal{R} of θ -approximate Ritz values of H has a cost $T_{\text{OptRitz}}(k, \theta, \delta)$. Then, using an efficient implicit QR algorithm (cf. Definition 2.1) each computation of $\text{iqr}(H, (z - r_i)^k)$ has a cost of $7kn^2$. By Lemma 4.5, we can produce a promising Ritz value in at most $12k \log kn^2 + \log k$ arithmetic operations. Then, in the event that the promising shift fails to reduce the potential the algorithm calls Exc, which takes $N_{\text{net}}(0.002 B^{-\frac{8 \log k + 4}{k-1}})$ arithmetic operations to specify the set \mathcal{S} of exceptional shifts. Some exceptional shift achieves potential reduction, and we pay $7kn^2$ operations for each one that we check. \square

4.3 Efficiently Finding a Promising Ritz Value

In this section we show how to efficiently find a promising Ritz value, in $O(n^2 k \log k)$ arithmetic operations. Note that it is trivial to find a $\kappa_\gamma(H)^{2/k}$ -promising Ritz value in $O(n^2 k^2)$ arithmetic operations simply by computing $\|e_n^*(H - r_i)^{-k/2}\|$ for $i = 1, \dots, k$ with k calls to $\text{iqr}(H, (z - r_i)^{k/2})$, choosing the maximizing index i , and appealing to Lemma 4.2. The content of Lemma 4.5 below that this can be done considerably more efficiently if we use a binary search type procedure. This improvement has nothing to do with the dynamical properties of our shifting strategy so readers uninterested in computational efficiency may skip this section.

Find

Input: Hessenberg H , a set $\mathcal{R} = \{r_1, \dots, r_k\}$ of θ -optimal Ritz values of H .

Output: A complex number $r \in \mathcal{R}$ with suitable properties that will be used in the main shift.

Requires: $\psi_k(H) > 0$

Ensures: r is α -promising for $\alpha = \kappa_V(H)^{\frac{4 \log k}{k}}$.

1. For $j = 1, \dots, \log k$
 - (a) Evenly partition $\mathcal{R} = \mathcal{R}_0 \sqcup \mathcal{R}_1$, and for $b = 0, 1$ set $p_{j,b} = \prod_{r \in \mathcal{R}_b} (z - r)$
 - (b) $\mathcal{R} \leftarrow \mathcal{R}_b$, where b maximizes $\|e_n^* p_{j,b}(H)^{-2^{j-1}}\|$
2. Output $\mathcal{R} = \{r\}$

Proof of Lemma 4.5 (Guarantees for Find) First, observe that $\|e_n^* q(H)\| \neq 0$ for every polynomial appearing in the definition of Find, since otherwise we would have $\psi_k(H) = 0$.

On the first step of the subroutine $p_{1,0} p_{1,1} = p$, the polynomial whose roots are the full set of approximate Ritz values, so

$$\begin{aligned} \max_b \|e_n^* p_{1,b}(H)^{-1}\| &\geq \frac{1}{\kappa_V(H)^2} \mathbb{E} \left[\frac{1}{2} \left(|p_{1,0}(Z_H)|^{-2} + |p_{1,1}(Z_H)|^{-2} \right) \right] && \text{Lemma 4.2} \\ &\geq \frac{1}{\kappa_V(H)^2} \mathbb{E}[|p(Z_H)|^{-1}] && \text{AM/GM.} \end{aligned}$$

On each subsequent step, we've arranged things so that $p_{j+1,0} p_{j+1,1} = p_{j,b}$, where b maximizes $\|e_n^* p_{j,b}(H)^{-2^{j-1}}\|$, and so by the same argument

$$\begin{aligned} &\max_b \|e_n^* p_{j+1,b}(H)^{-2^j}\|^2 \\ &\geq \frac{1}{\kappa_V(H)^2} \mathbb{E} \left[\frac{1}{2} \left(|p_{j+1,0}(Z_H)|^{-2^{j+1}} + |p_{j+1,1}(Z_H)|^{-2^{j+1}} \right) \right] && \text{Lemma 4.2} \\ &\geq \frac{1}{\kappa_V(H)^2} \mathbb{E} \left[|p_{j+1,0}(Z_H) p_{j+1,1}(Z_H)|^{-2^j} \right] && \text{AM/GM} \\ &\geq \frac{1}{\kappa_V(H)^4} \|e_n^* (p_{j+1,0}(H) p_{j+1,1}(H))^{-2^{j-1}}\| && \text{Lemma 4.2} \\ &= \frac{1}{\kappa_V(H)^4} \max_b \|e_n^* p_{j,b}(H)^{-2^{j-1}}\|. \end{aligned}$$

Paying a further $\kappa_V(H)^2$ on the final step to convert the norm into an expectation, we get

$$\mathbb{E} \left[|Z_H - r|^{-k} \right] \geq \frac{1}{\kappa_V(H)^{4 \log k}} \mathbb{E} \left[|p(Z_H)|^{-1} \right]$$

as promised.

For the runtime, we can compute every $\|e_n^* p_{j,b}(H)^{-2^{j-1}}\|$ by running an implicit QR step with the polynomials $p_{j,b}^{2^{j-1}}$, all of which have degree $k/2$. There are $2 \log k$ such computations throughout the subroutine, and each one requires $6kn^2$ arithmetic operations. Beyond that we need only compare the two norms on each of the $\log k$ steps. \square

Remark 4.8 (*Opportunism and Judicious Partitioning*) In practice, it may be beneficial to implement Find *opportunistically*, meaning that in each iteration one should check if the new set of Ritz values gives potential reduction (this can be combined with the computation of $\|e_n^* p_{j,b}(H)^{-2^{j-1}}\|$ and implemented with no extra cost). Moreover, note that Find does not specify a way to partition the set of Ritz values obtained after each iteration, and as can be seen from the above proof, the algorithm works regardless of the partitioning choices. It is conceivable that a judicious choice of the partitioning could be used to obtain further improvements.

4.4 Analysis of the Exceptional Shift

To conclude our analysis, it remains only to define the subroutine “Exc,” which produces a set \mathcal{S} of possible exceptional shifts in the event that an α -promising Ritz value does not achieve potential reduction. The main geometric intuition is captured in the case when H is normal and $\kappa_V(H) = 1$. Here, Find gives us a 1-promising Ritz value r and Lemma 4.6 with $t = 1/2$ tells us that if r does not achieve potential reduction, then Z_H has measure at least $\frac{1}{4}((1 - \gamma)/\theta)^{2k}$ on a disk of radius $R := 2^{1/k}\theta\psi_k(H)$.

For any $\epsilon > 0$, we can easily construct an $R\epsilon$ -net \mathcal{S} contained in this disk—i.e., a set with the property that every point in the disk is at least $R\epsilon$ -close to a point in \mathcal{S} —with $O(1/\epsilon)^2$ points. One can then find a point $s \in \mathcal{S}$ satisfying

$$\begin{aligned} \tau_{(z-s)^k}(H)^{-2k} &= \|e_n^*(H - s)^{-k}\|^2 \\ &= \mathbb{E}[|Z_H - s|^{-2k}] \\ &\geq \frac{\mathbb{P}[|Z_H - s| \leq \psi_k(H)]}{|\mathcal{S}|(R\epsilon)^{2k}} \\ &\approx \frac{1}{4} \left(\frac{(1 - \gamma)}{\theta} \right)^{2k} \frac{1}{R^{2k} \epsilon^{2k-2}}, \end{aligned}$$

where the first equality is by normality of H , and second inequality comes from choosing $s \in \mathcal{S}$ to maximize $|Z_H - s|^{-2k}$. Since $\psi_k(\text{iqr}(H, (z - s)^k)) \leq \tau_{(z-s)^k}(H)$, we can ensure potential reduction by setting $\epsilon \approx \frac{(1-\gamma)^2 R}{\theta \psi_k(H)} \approx ((1 - \gamma)/\theta)^2$.

When H is nonnormal, the chain of inequalities above hold only up to factors of $\kappa_V(H)$, and Find is only guaranteed to produce a $\kappa_V(H)^{4 \log k/k}$ -promising Ritz value. The necessary adjustments are addressed below in the implementation of Exc and the subsequent proof of its guarantees.

Exc

Input: Hessenberg H , a θ -approximate Ritz value r , a condition number bound B , promising parameter α

Output: A set $\mathcal{S} \subset \mathbb{C}$, which will be used as a set of exceptional shifts

Requires: $\kappa_V(H) \leq B$, r is α -promising, and $\psi_k(\text{iqr}(H, (z-r)^k)) \geq (1-\gamma)\psi_k(H)$

Ensures: For some $s \in \mathcal{S}$, $\psi_k(\text{iqr}(H, (z-s)^k)) \leq (1-\gamma)\psi_k(H)$

1. $R \leftarrow 2^{1/k} \theta \alpha B^{1/k} \psi_k(H)$
2. $\epsilon \leftarrow \left(\frac{(1-\gamma)^2}{(12B^4)^{1/k} \alpha^2 \theta^2} \right)^{\frac{k}{k-1}}$
3. $\mathcal{S} \leftarrow \epsilon R$ -net of $R\psi_k(H)$.

Proof of Lemma 4.7: Guarantees for Exc Instantiating $t = 1/2$ in equation (23), we find that for the setting of R in line (1) of Exc,

$$\mathbb{P}[|Z_H - r| \leq D(r, R)] \geq \frac{1}{4B^4} \left(\frac{(1-\gamma)}{\alpha\theta} \right)^{2k}.$$

Let \mathcal{S} be an ϵR -net of $D(r, R)$; it is routine that such a net has at most $(1+2/\epsilon)^2 \leq 9/\epsilon^2$ points. By Lemma 2.5, to show that some $s \in \mathcal{S}$ achieves potential reduction, it suffices to find one for which

$$\|e_n^*(H - s)^{-k}\|^2 \geq \frac{1}{(1-\gamma)^{2k} \psi_k(H)^{2k}}.$$

We thus compute

$$\begin{aligned} & \max_{s \in \mathcal{S}} \|e_n^*(H - s)^{-k}\|^2 \\ & \geq \frac{1}{\kappa_V(H)^2 |\mathcal{S}|} \sum_{s \in \mathcal{S}} \mathbb{E} \left[|Z_H - s|^{-2k} \right] \\ & \geq \frac{\epsilon^2}{9B^2} \mathbb{E} \left[\sum_{s \in \mathcal{S}} |Z_H - s|^{-2k} \cdot \mathbf{1}_{Z_H \in D(r, R)} \right] && \text{Fubini and } \kappa_V(H) \leq B \\ & \geq \frac{\epsilon^2}{9B^2} \mathbb{E} \left[\max_{s \in \mathcal{S}} |Z_H - s|^{-2k} \cdot \mathbf{1}_{Z_H \in D(r, R)} \right] \\ & \geq \frac{\epsilon^2}{9B^2} \mathbb{E} \left[\frac{\mathbf{1}_{Z_H \in D(r, R)}}{(\epsilon R)^{2k}} \right] && \mathcal{S} \text{ is an } \epsilon R \text{ - net} \\ & \geq \frac{\mathbb{P}[Z_H \in D(r, R)]}{9B^2 R^{2k} \epsilon^{2k-2}} \\ & \geq \frac{1}{(1-\gamma)^{2k} \psi(H)^{2k}} \end{aligned}$$

with the second to last line following from the fact that some $s \in \mathcal{S}$ is at least ϵR -close to Z_H whenever the latter is in $D(r, R)$, and the final inequality holding provided that

$$\epsilon \leq \left(\frac{\mathbb{P}[|Z_H - r| \leq R\psi_k(H)](1 - \gamma)^{2k}\psi_k(H)^{2k}}{9B^2R^{2k}} \right)^{\frac{1}{2k-2}}.$$

Expanding the probability and using the definition of R in line 1, it suffices to set ϵ smaller than

$$\begin{aligned} & \left(\frac{(1 - \gamma)^{2k}}{4B^4\alpha^{2k}\theta^{2k}} \cdot \frac{(1 - \gamma)^{2k}\psi_k(H)^{2k}}{9B^2} \cdot \frac{1}{4B^2\alpha^{2k}\theta^{2k}\psi_k(H)^{2k}} \right)^{\frac{1}{2k-2}} \\ & = \left(\frac{(1 - \gamma)^2}{(12B^4)^{1/k}\alpha^2\theta^2} \right)^{\frac{k}{k-1}}, \end{aligned}$$

which is the quantity appearing in line 2. Setting $\theta = 2$, $\gamma = 0.2$, and $\alpha = B^{4 \log k/k}$, and using $k \geq 2$, we obtain the expression appearing in $N_{\text{net}}(\cdot)$ in the statement of Lemma 4.7.

However, a more practical choice (and the one that we will use in the companion paper [6]) is an equilateral triangular lattice with spacing $\sqrt{3}\epsilon$, intersected with the $D(r, (1 + \epsilon)R)$. Such a construction is optimal as $\epsilon \rightarrow 0$, and can be used to give a better bound on $N_{\text{net}}(\epsilon)$ when ϵ is small. For instance, by adapting an argument of [2, Lemma 2.6] one can show that this choice of \mathcal{S} satisfies

$$N_{\text{net}}(\epsilon) \leq \frac{2\pi}{3\sqrt{3}}(1 + 1/\epsilon)^2 + \frac{4\sqrt{2}}{\sqrt{3}}(1 + 1/\epsilon) + 1.$$

In the normal case, when $B = \alpha = \theta = 1$, $k = 4$, and $\gamma = 0.2$, the above bound gives

$$|\mathcal{S}| \leq N_{\text{net}} \left(\left(\frac{0.8^2}{12^{1/4}} \right)^{4/3} \right) \leq 49.9.$$

□

Acknowledgements We thank Jim Demmel, Daniel Kressner, and Cleve Moler for helpful conversations and references. We thank the anonymous referees for a thorough reading and feedback which greatly improved the paper.

References

1. Kensuke Aishima, Takayasu Matsuo, Kazuo Murota, and Masaaki Sugihara. A Wilkinson-like multishift QR algorithm for symmetric eigenvalue problems and its global convergence. *Journal of Computational and Applied Mathematics*, 236(15):3556–3560, 2012.

2. Diego Armentano, Carlos Beltrán, Peter Bürgisser, Felipe Cucker, and Michael Shub. A stable, polynomial-time algorithm for the eigenpair problem. *Journal of the European Mathematical Society*, 20(6):1375–1437, 2018.
3. Zhaojun Bai and James Demmel. On a block implementation of Hessenberg multishift QR iteration. *International Journal of High Speed Computing*, 1(01):97–112, 1989.
4. Jess Banks, Jorge Garza-Vargas, Archit Kulkarni, and Nikhil Srivastava. Overlaps, eigenvalue gaps, and pseudospectrum under real Ginibre and absolutely continuous perturbations. *arXiv preprint arXiv:2005.08930*, 2020.
5. Jess Banks, Jorge Garza-Vargas, Archit Kulkarni, and Nikhil Srivastava. Pseudospectral shattering, the sign function, and diagonalization in nearly matrix multiplication time. *Foundations of Computational Mathematics*, pages 1–89, 2022.
6. Jess Banks, Jorge Garza-Vargas, and Nikhil Srivastava. Global linear convergence of Hessenberg shifted QR II: Numerical stability. *arXiv preprint*, 2022.
7. Jess Banks, Jorge Garza-Vargas, and Nikhil Srivastava. Global linear convergence of Hessenberg shifted QR III: Approximate Ritz values via shifted inverse iteration. *arXiv preprint*, 2022.
8. Jess Banks, Archit Kulkarni, Satyaki Mukherjee, and Nikhil Srivastava. Gaussian regularization of the pseudospectrum and Davies' conjecture. *Communications on Pure and Applied Mathematics*, 74(10):2114–2131, 2021.
9. Steve Batterson. Convergence of the shifted QR algorithm on 3×3 normal matrices. *Numerische Mathematik*, 58(1):341–352, 1990.
10. Steve Batterson. Convergence of the Francis shifted QR algorithm on normal matrices. *Linear algebra and its applications*, 207:181–195, 1994.
11. Steve Batterson. Dynamical analysis of numerical systems. *Numerical linear algebra with applications*, 2(3):297–310, 1995.
12. Steve Batterson and David Day. Linear convergence in the shifted QR algorithm. *Mathematics of Computation*, 59(199):141–151, 1992.
13. Steve Batterson and John Smillie. The dynamics of Rayleigh quotient iteration. *SIAM Journal on Numerical Analysis*, 26(3):624–636, 1989.
14. Steve Batterson and John Smillie. Rayleigh quotient iteration for nonsymmetric matrices. *Mathematics of Computation*, 55(191):169–178, 1990.
15. Karen Braman, Ralph Byers, and Roy Mathias. The multishift QR algorithm. Part I: Maintaining well-focused shifts and level 3 performance. *SIAM Journal on Matrix Analysis and Applications*, 23(4):929–947, 2002.
16. Karen Braman, Ralph Byers, and Roy Mathias. The multishift QR algorithm. Part II: Aggressive early deflation. *SIAM Journal on Matrix Analysis and Applications*, 23(4):948–973, 2002.
17. Hendrik Jan Buurema. *A geometric proof of convergence for the QR method*, volume 62. 1958.
18. Ralph Byers. Lapack 3.1 xHSEQR: Tuning and implementation notes on the small bulge multi-shift QR algorithm with aggressive early deflation. Technical report, LAPACK Working Note 187, 2007.
19. Moody T Chu. Linear algebra algorithms as dynamical systems. *Acta Numerica*, 17:–86, 2008.
20. E. Brian Davies. Approximate diagonalization. *SIAM Journal on Matrix Analysis and Applications*, 29(4):1051–1064, 2008.
21. David Day. How the QR algorithm fails to converge and how to fix it. 1996.
22. Theodorus J Dekker and Joseph F. Traub. The shifted QR algorithm for Hermitian matrices. *Linear Algebra Applications*, 4:137–154, 1971.
23. James W Demmel. *Applied numerical linear algebra*. SIAM, 1997.
24. Jack Dongarra and Francis Sullivan. Guest editors' introduction: The top 10 algorithms. *IEEE Computer Architecture Letters*, 2(01):22–23, 2000.
25. Patricia J Eberlein and C. P. Huang. Global convergence of the QR algorithm for unitary matrices with some results for normal matrices. *SIAM Journal on Numerical Analysis*, 12(1):97–104, 1975.
26. Jiang Erxiong. A note on the double-shift QL algorithm. *Linear algebra and its applications*, 171:121–132, 1992.
27. John GF Francis. The QR transformation a unitary analogue to the LR transformation-Part 1. *The Computer Journal*, 4(3):265–271, 1961.
28. John GF Francis. The QR transformation-Part 2. *The Computer Journal*, 4(4):332–345, 1962.
29. Gene Golub and Frank Uhlig. The QR algorithm: 50 years later its genesis by John Francis and Vera Kublanovskaya and subsequent developments. *IMA Journal of Numerical Analysis*, 29(3):467–485, 2009.

30. Gene H Golub and Charles F Van Loan. Matrix computations. Johns Hopkins studies in the mathematical sciences, 1996.
31. Nicholas J Higham. *Accuracy and stability of numerical algorithms*. SIAM, 2002.
32. Nicholas J Higham, Mark R Dennis, Paul Glendinning, Paul A Martin, Fadil Santosa, and Jared Tanner. *The Princeton companion to applied mathematics*. Princeton University Press Princeton, NJ, USA:, 2015.
33. Walter Hoffmann and Beresford N Parlett. A new proof of global convergence for the tridiagonal QL algorithm. *SIAM Journal on Numerical Analysis*, 15(5):929–937, 1978.
34. Vishesh Jain, Ashwin Sah, and Mehtaab Sawhney. On the real Davies’ conjecture. *arXiv preprint arXiv:2005.08908*, 2020.
35. Daniel Kressner. *Personal communication*, 2021.
36. Vera N Kublanovskaya. On some algorithms for the solution of the complete eigenvalue problem. *USSR Computational Mathematics and Mathematical Physics*, 1(3):637–657, 1962.
37. Ricardo S Leite, Nicolau C Saldanha, and Carlos Tomei. Dynamics of the symmetric eigenvalue problem with shift strategies. *International Mathematics Research Notices*, 2013(19):4382–4412, 2013.
38. Cleve Moler. Variants of the QR algorithm. *Cleve’s Corner, Mathworks Technical Articles*, 2014.
39. Cleve B Moler. Three research problems in numerical linear algebra. *AMS Proceedings of Symposia in Applied Math*, vol. 22, pages 1–18, 1978.
40. Alexander M Ostrowski. On the convergence of the Rayleigh quotient iteration for the computation of the characteristic roots and vectors. I. *Archive for Rational Mechanics and Analysis*, 1(1):233–241, 1957.
41. Victor Y Pan. Univariate polynomials: nearly optimal algorithms for numerical factorization and root-finding. *Journal of Symbolic Computation*, 33(5):701–733, 2002.
42. Beresford Parlett. Singular and invariant matrices under the QR transformation. *Mathematics of Computation*, 20(96):611–615, 1966.
43. Beresford Parlett. Normal Hessenberg and moment matrices. *Linear Algebra and its Applications*, 6:37–43, 1973.
44. Beresford N Parlett. The Rayleigh quotient iteration and some generalizations for nonnormal matrices. *Mathematics of Computation*, 28(127):679–693, 1974.
45. Beresford N Parlett. The QR algorithm. *Computing in Science & Engineering*, 2(1):38–42, 2000.
46. Beresford N Parlett and William Kahan. On the convergence of a practical QR algorithm. In *IFIP Congress (I)*, pages 114–118, 1968.
47. Beresford N Parlett and Jian Le. Forward instability of tridiagonal QR. *SIAM Journal on Matrix Analysis and Applications*, 14(1):279–316, 1993.
48. Michael Reed and Barry Simon. *I: Functional analysis*, volume 1. Gulf Professional Publishing, 1980.
49. Edward B Saff and Vilmos Totik. *Logarithmic potentials with external fields*, volume 316. Springer Science & Business Media, 2013.
50. Steve Smale. Complexity theory and numerical analysis. *Acta numerica*, 6:523–551, 1997.
51. Daniel A Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM (JACM)*, 51(3):385–463, 2004.
52. Lloyd N Trefethen and David Bau III. *Numerical linear algebra*, volume 50. SIAM, 1997.
53. Tai-Lin Wang. Convergence of the tridiagonal QR algorithm. *Linear algebra and its applications*, 322(1-3):1–17, 2001.
54. Tai-Lin Wang and William Gragg. Convergence of the shifted QR algorithm for unitary Hessenberg matrices. *Mathematics of computation*, 71(240):1473–1496, 2002.
55. Tai-Lin Wang and William Gragg. Convergence of the unitary QR algorithm with a unimodular Wilkinson shift. *Mathematics of computation*, 72(241):375–385, 2003.
56. David S Watkins. Understanding the QR algorithm. *SIAM review*, 24(4):427–440, 1982.
57. David S Watkins. Forward stability and transmission of shifts in the QR algorithm. *SIAM Journal on Matrix Analysis and Applications*, 16(2):469–487, 1995.
58. David S Watkins. *The matrix eigenvalue problem: GR and Krylov subspace methods*. SIAM, 2007.
59. David S Watkins. The QR algorithm revisited. *SIAM review*, 50(1):133–145, 2008.
60. James H Wilkinson. Global convergence of tridiagonal QR algorithm with origin shifts. *Linear Algebra and its Applications*, 1(3):409–420, 1968.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.