

ORBITAL MINIMIZATION METHOD WITH ℓ^1 REGULARIZATION

JIANFENG LU AND KYLE THICKE

ABSTRACT. We consider a modification of the OMM energy functional which contains an ℓ^1 penalty term in order to find a sparse representation of the low-lying eigenspace of self-adjoint operators. We analyze the local minima of the modified functional as well as the convergence of the modified functional to the original functional. Algorithms combining soft thresholding with gradient descent are proposed for minimizing this new functional. Numerical tests validate our approach. As an added bonus, we also prove the unanticipated and remarkable property that every local minimum the OMM functional without the ℓ^1 term is also a global minimum.

1. INTRODUCTION

This paper considers solving for a sparse representation of the low-lying eigenspace of self-adjoint operators. Given a Hermitian matrix H , the low-lying eigenspace is defined as the linear combination of the first m eigenvectors of H . Our aim is to find a sparse representation by solving a variational problem for the eigenspace which at the same time favors sparsity.

Such sparse representations have applications in many scenarios, including constructing a localized numerical basis for PDE problems and sparse PCA in statistics. The main motivation of the current work comes from electronic structure calculations, for which the sparse representations, given by (approximate) linear combinations of the low-lying eigenfunctions, are known as the Wannier functions [7, 21].

Motivated by the recent works [17, 9, 8] that exploit the ℓ^1 penalty to enhance sparsity (see also [3, 12] for alternative localization strategies), in this work we consider minimizing the variational model

$$(1) \quad E_\mu(X) = \text{tr}[(2I - X^*X)X^*HX] + \mu \|X\|_1,$$

where $X \in \mathbb{C}^{N \times m}$ and $\|\cdot\|_1$ is the entry-wise ℓ^1 norm. Here, the energy functional without the ℓ^1 penalty

$$(2) \quad E_0(X) = \text{tr}[(2I - X^*X)X^*HX],$$

is used in the orbital minimization method (OMM), developed in the context of linear scaling algorithms for electronic structure [14, 13, 15, 16] to alleviate the orthogonality constraint $X^*X = I$. Hence, we will refer to E_0 as the OMM functional and refer to E_μ as the OMM functional with an ℓ^1 penalty.

Since the orthogonality constraint is lifted, E_μ can be minimized with unconstrained minimization algorithms. This allows the algorithm to be significantly simpler than that for the trace minimization with an ℓ^1 penalty proposed in [17] based on operator splitting [10]:

$$(3) \quad E_\mu^{\text{trace}}(X) = \text{tr}[X^*HX] + \mu \|X\|, \quad \text{s.t. } X^*X = I.$$

Date: May 4, 2016.

This work is partially supported by the National Science Foundation under grants DMS-1454939 and ACI-1450280. We thank Fabiano Corsetti, Haizhao Yang and Wotao Yin for helpful discussions.

Note that it is possible to lift the orthogonality constraint by convexification, as in [9, 8], which leads to localized density matrix minimization. However, the density matrix $P \in \mathbb{C}^{N \times N}$, which is the projection operator onto the low-lying eigenspace, contains many more degrees of freedom than X when $N \gg m$. Hence, it might be favorable to consider the non-convex functional E_μ with fewer degrees of freedom.

For the OMM functional without the ℓ^1 penalty, it is well known [14, 18] that the global minimizers of the OMM functional E_0 correspond to a basis for the low-lying eigenspace of H , if H is negative definite. Somewhat surprisingly, it turns out that we can further show that E_0 has no local minima in the sense that every local minimum is a global minimum, and hence a representation of the low-lying eigenspace, as shown in Theorem 2. In particular, when minimizing E_0 using e.g., the conjugate gradient method, we will not be trapped at local minima.

With the ℓ^1 term, the minimizer of E_μ no longer corresponds to the exact low-lying eigenspace; however, we show that it gives a good approximation as $\mu \rightarrow 0$. We will further analyze the approximate orthogonality of the minimizer and approximation to the density matrix.

The OMM functional has been used to develop linear scaling electronic structure algorithms (so that computational cost scales linearly with the number of electrons). See e.g., the review paper [5]. The conventional strategy is to simply truncate the domain of the problem by minimizing E_0 over a subspace of matrices in $\mathbb{C}^{N \times m}$ which have a particular predefined support [14, 13] with only $O(m)$ degrees of freedom. However, this truncation is known to create many local minima in the problem, which trap the minimization. While efforts to alleviate the local minima issue have been undertaken [6, 4, 20], it still remains a challenge for practical applications. It turns out that minimizing E_μ naturally leads to a soft thresholding optimization algorithm, thanks to the ℓ^1 penalty term. In our numerical tests, we find that local minima are not encountered in practice unless the penalty parameter μ is large (details in Section 4.3). Our method also has the additional advantage that the support of the solution does not have to be chosen ahead of time.

We also note that even in the context of a cubic scaling implementation, the OMM algorithm still has an advantage over direct eigensolvers in terms of scalability in parallel implementations, as recently demonstrated in [2]. The inclusion of the ℓ^1 penalty controls the sparsity of the solution, which allows for sparse matrix arithmetic and will hopefully enable a natural transition between the cubic scaling and reduced scaling algorithms.

The remainder of the paper is structured as follows. In Section 2, we introduce the different energy functionals we will be working with. We perform analysis to examine the local minima of the functionals and to examine the convergence of the modified functional to the original functional. In Section 3, we present algorithms for numerically minimizing our new functional. Numerical tests are performed in Section 4 to validate the theoretical results from Section 2 as well as to examine the performance of the algorithms proposed in Section 3.

2. ANALYSIS

2.1. Original OMM energy functional. When using Kohn–Sham density functional theory, one encounters the problem of calculating the low-lying eigenspace of a matrix H . This can be rephrased as the following minimization problem. Minimize the following functional over the set of matrices $Y \in \mathbb{C}^{N \times m}$ satisfying the orthonormality constraint $Y^* Y = I$.

$$(4) \quad E_\perp(Y) = \text{tr}[Y^* H Y].$$

The orthonormality constraint can be difficult to deal with, but we can relax the constraint by replacing Y with $XS^{-1/2}$ where $S = X^*X$. If X has full rank, then $XS^{-1/2}$ satisfies the orthonormality constraint. By making this substitution, we obtain the functional

$$(5) \quad E_{S^{-1}}(X) = \text{tr}[S^{-1}X^*HX].$$

Minimizing (5) is clearly equivalent to minimizing (4), but we have relaxed the restrictions on the set we are minimizing over. Instead of minimizing only over the set of matrices with orthonormal columns, we can now minimize over the set of all full rank matrices in $\mathbb{C}^{N \times m}$.

The idea behind the OMM functional [14, 13] is to avoid computing the inverse of S , by replacing S^{-1} by an approximation obtained by a Neumann expansion around I , the identity matrix. In particular, we can replace S^{-1} with a matrix of the following form

$$(6) \quad Q_{\mathcal{N}} = \sum_{k=0}^{\mathcal{N}} (I - S)^k.$$

Then instead of minimizing $E_{S^{-1}}$, we can consider minimizing the functional

$$(7) \quad E_{Q_{\mathcal{N}}}(X) = \text{tr}[Q_{\mathcal{N}}X^*HX],$$

where \mathcal{N} is an odd integer (if \mathcal{N} is even, the functional is not bounded from below). It has been shown that this alternate functional has very nice properties [14, 13] which we summarize in the following theorem.

Theorem 1. *Let $H \in \mathbb{C}^{N \times N}$ be Hermitian with eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N$. Then the global minima of $E_{S^{-1}}$ is the set of matrices $X \in \mathbb{C}^{N \times m}$ whose column space is a span of eigenvectors corresponding to the m most negative eigenvalues of X .*

Suppose further that H is negative definite (we can always shift $H \rightarrow H - \eta I$ if necessary). Let \mathcal{N} be an odd positive integer. Then, the functional $E_{Q_{\mathcal{N}}}$ has a global minimum, and in fact, $\min E_{Q_{\mathcal{N}}} = \min E_{S^{-1}}$. Additionally, X is a minimizer of $E_{Q_{\mathcal{N}}}$ if and only if X is a minimizer of $E_{S^{-1}}$ and the columns of X are orthonormal.

Note that the ‘‘eigenvectors corresponding to the m most negative eigenvalues of H ’’ may not be unique since it could be that $\lambda_m = \lambda_{m+1}$. When we use this phrase, we are referring to any possible choice of eigenvectors, when degeneracy occurs.

The major conclusion of the theorem is that we can find exact minimizers of $E_{S^{-1}}$ by minimizing the simpler functional $E_{Q_{\mathcal{N}}}$ which does not require the computation of S^{-1} . This both reduces the computational cost and allows us to minimize over all matrices rather than just those of full rank. In this paper, we will choose $\mathcal{N} = 1$. In particular, we consider the problem of minimizing the functional we call E_0 , defined by

$$(8) \quad E_0(X) = \text{tr}[(2I - X^*X)X^*HX],$$

where $X \in \mathbb{C}^{N \times m}$ and $H \in \mathbb{C}^{N \times N}$ is a Hermitian negative definite matrix with eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_N < 0$.

Our first main result is the following theorem. The result is quite unexpected as E_0 is clearly non-convex, while the theorem states that local minima of E_0 are also global minima.

Theorem 2. *Any local minimum of E_0 is also a global minimum.*

Before proving the theorem, we note a very practical consequence of it. Many numerical optimization algorithms can guarantee convergence to a critical point of the objective function. If we use such an algorithm, then after the algorithm converges, we can perturb the solution.

Theorem 2 tells us that if starting the algorithm from this perturbed point leads to convergence to the same value of E_0 , then that point is likely a global minimum.

The proof of Theorem 2 will follow from the following lemmas. The first two lemmas we state are elementary. Then the next two lemmas will provide the crucial tools needed to prove the theorem. Lemma 2.3 provides a characterization of the critical points of E_0 , and Lemma 2.4 allows us to analyze the behavior of E_0 near the critical points. By using the analysis in Lemma 2.4, we can examine all the critical points which have been characterized in Lemma 2.3 to show that all the critical points which are local minima are actually global minima.

Lemma 2.1. *If $A \in \mathbb{C}^{n \times n}$ is Hermitian and V is a k -dimensional invariant subspace of A , then there exist k eigenvectors $\{\xi_i\}_{i=1}^k$ of A such that $V = \text{span}(\{\xi_i\}_{i=1}^k)$.*

Proof. Let $\{b_i\}_{i=1}^k$ be an orthonormal basis for V . Define $B = [b_1, \dots, b_k]$. Then $A|_V$ with respect to the basis $\{b_i\}_{i=1}^k$ is B^*AB , which is Hermitian. Therefore, $A|_V$ is diagonalizable with eigenvectors $\{\eta_i\}_{i=1}^k \subset \mathbb{C}^k$. This means $B^*AB\eta_i = \lambda_i\eta_i$. Multiplying on the left by B and noting that $BB^* = I$ on V , we see that $\{B\eta_i\}_{i=1}^k$ are eigenvectors of A . It is easy to see that $\{B\eta_i\}_{i=1}^k$ are independent and therefore form a basis for V . \square

Lemma 2.2. *E_0 is invariant under right multiplication by elements of $U(m)$. That is, $E_0(XG) = E_0(X)$ for all $X \in \mathbb{C}^{N \times m}$ and $G \in U(m)$.*

Proof. Use basic properties of the trace. \square

Before continuing, we note that since X^*X is Hermitian, there exists $G \in U(m)$ such that $X^*X = G\Lambda G^*$ where Λ is diagonal. We will often combine this fact with Lemma 2.2 to assume without loss of generality that X^*X is diagonal. We will now prove the above assertions concerning E_0 by following a natural progression through the following lemmas. The next lemma is interesting in its own right as it provides a characterization of the critical points of E_0 . Its proof proceeds in two steps. First, we show that the eigenvalues of X^*X are in $\{0, 1\}$. Then, we can use this fact to show that the column space of X , $\text{Col}(X)$, is the span of some set of eigenvectors of H .

Lemma 2.3. *If $\nabla E_0(X) = 0$, then all eigenvalues of X^*X are in $\{0, 1\}$. Additionally, the $\text{Col}(X)$ is the span of some set of eigenvectors of H .*

Proof. Note that if $G \in U(m)$, then $\nabla E_0(XG) = \nabla E_0(X)G$. So without loss of generality, assume that X^*X is diagonal. By assumption we have,

$$\begin{aligned} \nabla E_0(X) &= -2XX^*HX + 4HX - 2HXX^*X \\ (9) \quad &= -2(XX^*H + HXX^* - 2H)X = 0. \end{aligned}$$

Denote $X = [x_1, x_2, \dots, x_m]$, and let $T = XX^*H + HXX^* - 2H$. Note that $XX^* = \sum x_i x_i^*$. By noting that the columns of X are orthogonal (since X^*X is diagonal), we have

$$(10) \quad XX^*x_i = x_i x_i^* x_i = \|x_i\|^2 x_i,$$

$$(11) \quad x_i^* XX^* = x_i^* x_i x_i^* = \|x_i\|^2 x_i^*.$$

Now we use the fact that (9) implies $x_i^* T x_i = 0$.

$$\begin{aligned} 0 &= x_i^* (XX^*H + HXX^* - 2H)x_i \\ (12) \quad &= 2(\|x_i\|^2 - 1)x_i^* H x_i. \end{aligned}$$

H is negative definite, so if $x_i \neq \vec{0}$, then $x_i^* H x_i < 0$. This implies that either $x_i = \vec{0}$ or $\|x_i\| = 1$. Since the diagonal entries of $X^* X$ are $\|x_i\|^2$, this implies that all eigenvalues of $X^* X$ are contained in the set $\{0, 1\}$.

Now we may assume without loss of generality that $X^* X$ is diagonal with diagonal entries consisting only of 0s and 1s. Equation (9) implies that for each column x_i of X ,

$$\begin{aligned} \vec{0} &= (X X^* H + H X X^* - 2H)x_i \\ &= X X^* H x_i + H x_i - 2H x_i \\ (13) \quad &= (X X^* - I)H x_i. \end{aligned}$$

But $\text{Null}(X X^* - I) = \text{Col}(X)$, so $H x_i \in \text{Col}(X)$ for all i . This implies that $\text{Col}(X)$ is an invariant subspace of H . So, by Lemma 2.1, $\text{Col}(X)$ is the span of some set of eigenvectors of H . \square

Now we consider the perturbation of the energy of E_0 around a critical point. The analysis is similar to, but more general than that in [18]. The lemma will be used to prove that all local minima of E_0 are global minima.

Lemma 2.4. *Let $\{y_i\}_{i=1}^N$ be a complete orthonormal set of eigenvectors of H corresponding respectively to the (not necessarily ordered) eigenvalues $\{\lambda_i\}_{i=1}^N$ of H . Let $Z = [z_1, z_2, \dots, z_m]$ where z_i is either y_i or $\vec{0}$. Let $X = [x_1, x_2, \dots, x_m]$ where $x_i = z_i + \sum_{j=1}^N c_j^{(i)} y_j$. Then,*

$$\begin{aligned} E_0(X) - E_0(Z) &= \sum_{i=1}^m \sum_{k=m+1}^N |c_k^{(i)}|^2 \left(2\lambda_k - (\lambda_k + \lambda_i)\chi_i \right) + 2 \sum_{i=1}^m \left[(1 - \chi_i) |c_i^{(i)}|^2 - 2(\text{Re}(c_i^{(i)}))^2 \chi_i \right] \lambda_i \\ &\quad - \sum_{i=1}^m \sum_{\substack{j=1 \\ i \neq j}}^m \left[c_i^{(j)} c_j^{(i)} \lambda_i \chi_i \chi_j + c_i^{(j)*} c_j^{(i)*} \lambda_j \chi_i \chi_j + |c_i^{(j)}|^2 \lambda_i \chi_i + |c_j^{(i)}|^2 \lambda_j \chi_j \right. \\ &\quad \left. - |c_j^{(i)}|^2 \left(2\lambda_j - (\lambda_j + \lambda_i)\chi_i \right) \right] \\ (14) \quad &+ \mathcal{O}(\|X - Z\|_F^3), \end{aligned}$$

where $\chi_i = 0$ if $z_i = \vec{0}$ and $\chi_i = 1$ if $z_i = y_i$.

Proof. The calculation is tedious but straightforward. \square

Using the above lemmas, we may now prove Theorem 2 as promised.

Proof of Theorem 2. We first note that the global minima of E_0 are the X whose columns span the same space as the eigenvectors corresponding to the m most negative eigenvalues (counting geometric multiplicity). Due to the fact that eigenvalues can have more than one independent eigenvector, this m -dimensional ‘‘minimal subspace’’ may not be unique.

If X_0 is a local minimum, then Lemma 2.3 implies that the column space of X_0 is spanned by some set $\{y_i\}_{i=1}^r$ of orthonormal eigenvectors of H . Without loss of generality, we can assume that $X_0 = [y_1, y_2, \dots, y_r, \vec{0}, \dots, \vec{0}]$. We can extend the set of y_i 's to a complete orthonormal set of eigenvectors of H , $\{y_i\}_{i=1}^N$. Then we can apply Lemma 2.4 with $Z = X_0$.

Our last step is to show that if the columns of X_0 are not eigenvectors corresponding to the m most negative eigenvalues of H , then X_0 is not a local minimum. First, if X_0 contains a column which is $\vec{0}$, then it is easy to see that there must exist $C_1 > 0$ such that for $0 < |c_m^{(m)}| < C_1$, $E_0(X) < E_0(X_0)$, where X is defined as in Lemma 2.4 with all $c_i^{(j)} = 0$ except for $c_m^{(m)}$. So we may assume that X_0 does not contain any $\vec{0}$ columns. Next, if column s of X_0 is an eigenvector which does

not correspond to one of the m most negative eigenvalues, then an eigenvector y_t ($t > m$) that does correspond to one of the m most negative eigenvectors is not a column of X_0 . Define X as in Lemma 2.4 with all $c_i^{(j)} = 0$ except for $c_t^{(s)}$. Then it is easy to see that there exists $C_2 > 0$ such that for $0 < |c_t^{(s)}| < C_2$, $E_0(X) < E_0(X_0)$. This proves that X_0 is a local minimum only if its columns are eigenvectors corresponding to the m most negative eigenvalues of H . Therefore, X_0 is a global minimum. \square

We remark that we can say a little more to classify the critical points of E_0 . It is apparent from (51) that the only local maximum of E_0 is at the origin. Therefore, all non-origin critical points are either saddle points or global minima.

2.2. OMM energy functional with ℓ^1 penalty. The major goal for the rest of the paper is to introduce and analyze a modification of the OMM energy functional which favors sparsity in its solution. In order to favor sparsity, we add an ℓ^1 penalty term to the OMM energy functional. In particular, the new functional we consider is

$$(15) \quad E_\mu(X) = E_0(X) + \mu \|X\|_1,$$

where $\|X\|_1 = \sum_i \sum_j |X_{ij}|$ is the entrywise ℓ^1 norm (viewing X as a vector). By minimizing (15), we hope to find solutions which are both sparse and “near” the global minima of E_0 . In particular, as $\mu \rightarrow 0^+$, we would hope that the minima of E_μ approach the minima of E_0 . In Lemma 2.5, we will show that this is true. Additionally, in Theorem 3, we quantify the rate at which the minima of E_μ approach the minima of E_0 as $\mu \rightarrow 0^+$. Before stating these results, we define some notation that we will use throughout.

$$(16) \quad S_\mu = \arg \min E_\mu(X),$$

$$(17) \quad d(A, B) = \sup_{a \in A} \inf_{b \in B} \|a - b\|_F.$$

Now we prove that the set of minimizers of E_μ converges to the set of minimizers of E_0 as $\mu \rightarrow 0^+$. Even more, S_μ converges to the elements of S_0 which have minimal ℓ^1 norm.

Lemma 2.5. (Convergence as $\mu \rightarrow 0^+$) *Let $M = \arg \min_{X \in S_0} \|X\|_1$. Then, $d(S_\mu, M) \rightarrow 0$ as $\mu \rightarrow 0^+$.*

Proof. First, we note that by the definition of S_μ ($\mu > 0$), if $X \in S_\mu$ and $Y \in S_0$, then

$$(18) \quad E_0(X) + \mu \|X\|_1 = E_\mu(X) \leq E_\mu(Y) = \min E_0 + \mu \|Y\|_1$$

This implies that $\|X\|_1 \leq \|Y\|_1$. So, for all $\mu > 0$,

$$(19) \quad \sup_{X \in S_\mu} \|X\|_1 \leq \inf_{Y \in S_0} \|Y\|_1.$$

Now, let $\{\mu_k\}$ be a sequence that converges down to 0. Let $X_k \in S_{\mu_k}$ and $Y \in S_0$. By definition of S_{μ_k} ,

$$(20) \quad E_0(X_k) + \mu_k \|X_k\|_1 = E_{\mu_k}(X_k) \leq E_{\mu_k}(Y) = \min E_0 + \mu_k \|Y\|_1$$

which implies

$$(21) \quad E_0(X_k) - \min E_0 \leq \mu_k (\|Y\|_1 - \|X_k\|_1) \leq \mu_k \|Y\|_1$$

Since this holds for all $k \in \mathbb{N}$ and $\mu_k \rightarrow 0$, we conclude that $E_0(X_k) \rightarrow \min E_0$. Now suppose by way of contradiction that there exists $\varepsilon > 0$ and a subsequence $\{X_{k_l}\}$ such that $d(X_{k_l}, S_0) > \varepsilon$ for all l . Eq. (19) implies that there exists a convergent subsequence $\{X_{k_{l_m}}\}$. And by continuity of

E_0 , we get our contradiction: $X_{k_l m} \rightarrow X^* \in S_0$. Since μ_k and X_k were arbitrary, we have shown that $d(S_\mu, S_0) \rightarrow 0$ as $\mu \rightarrow 0^+$. Now, by (19), it is clear that $d(S_\mu, M) \rightarrow 0$ as $\mu \rightarrow 0^+$. \square

Now that we have proven convergence, we further state and prove the convergence rate.

Theorem 3. (Convergence rate as $\mu \rightarrow 0^+$) *Let $\alpha \in (0, 1)$. Suppose $\lambda_m < \lambda_{m+1}$. Then there exists $\mu_0(\alpha) > 0$ such that for $0 < \mu < \mu_0(\alpha)$,*

$$(22) \quad d(S_\mu, S_0) \leq \frac{\mu \sqrt{Nm}}{\alpha(\lambda_{m+1} - \lambda_m)}.$$

Proof. Without loss of generality, assume that the eigenvectors of H associated to $\lambda_1, \dots, \lambda_N$ are the standard basis vectors, $\vec{e}_1, \dots, \vec{e}_N$, respectively (a change of basis could always be used to bring the problem into this form). Then, Theorem 1 implies that the global minimizers of E_0 are given by the set

$$(23) \quad \operatorname{argmin} E_0 = \left\{ \begin{bmatrix} G \\ 0 \end{bmatrix} \in \mathbb{C}^{N \times m} : G \in U(m) \right\}.$$

Let $X \in \mathbb{C}^{N \times m}$. Since the set $\operatorname{argmin} E_0$ is closed, the set $\operatorname{argmin} \{\|X - G\|_F : G \in \operatorname{argmin} E_0\}$ is nonempty. By Lemma 2.2, we can rotate X without changing $E_0(X)$ such that

$$\tilde{I} = \begin{bmatrix} I_m \\ 0 \end{bmatrix} \in \operatorname{argmin} \{\|X - G\|_F : G \in \operatorname{argmin} E_0\}.$$

Next, we wish to show that for such X , we have $X_{ij} = X_{ji}$ for $i < j \leq m$. For a given i, j, θ , define $G^{(i,j)}(\theta) \in \mathbb{C}^{N \times m}$ such that

$$(24) \quad G_{ii}^{(i,j)}(\theta) = G_{jj}^{(i,j)}(\theta) = \cos(\theta),$$

$$(25) \quad G_{ij}^{(i,j)}(\theta) = -G_{ji}^{(i,j)}(\theta) = \sin(\theta),$$

$$(26) \quad G_{kl}^{(i,j)}(\theta) = \tilde{I}_{kl}, \quad \text{for } (k, l) \notin \{(i, i), (i, j), (j, i), (j, j)\}.$$

Note that $G^{(i,j)}(\theta) \in \operatorname{argmin} E_0$ for all θ , and $G^{(i,j)}(0) = \tilde{I}$. It easy to show that

$$(27) \quad \frac{d}{d\theta} \left\| X - G^{(i,j)}(\theta) \right\|_F^2 \Big|_{\theta=0} = 2 \operatorname{Re}(X_{ji} - X_{ij}),$$

$$(28) \quad \frac{d}{d\theta} \left\| X - iG^{(i,j)}(\theta) \right\|_F^2 \Big|_{\theta=0} = 2 \operatorname{Im}(X_{ji} - X_{ij}).$$

But these must equal 0 since $\tilde{I} \in \operatorname{argmin} \{\|X - G\|_F : G \in \operatorname{argmin} E_0\}$. Therefore, $X_{ij} = X_{ji}$. Next, we show that $\operatorname{Im}(X_{jj}) = 0$ for $1 \leq j \leq m$. Define $G^{(j,j)}$ as \tilde{I} except with $G_{jj}^{(j,j)} = 1 + i \operatorname{Im}(X_{jj})$. It is clear that if $\operatorname{Im}(X_{jj}) \neq 0$, then $\|X - G^{(j,j)}\|_F < \|X - \tilde{I}\|_F$ which contradicts the assumption that $\tilde{I} \in \operatorname{argmin}_{G \in \operatorname{argmin} E_0} \|X - G\|_F$. Therefore, $\operatorname{Im}(X_{jj}) = 0$ for all j .

Using these facts along with Lemma 2.4, we obtain

$$(29) \quad E_0(X) - \min E_0 = \sum_{i=1}^m \sum_{k=m+1}^N (\lambda_k - \lambda_i) |X_{ik} - \tilde{I}_{ik}|^2 + 4 \sum_{i=1}^m |\lambda_i| (\operatorname{Re}(X_{ii} - \tilde{I}_{ii}))^2 + \sum_{i < j} (|\lambda_i| + |\lambda_j|) |(X_{ji} - \tilde{I}_{ji}) + (X_{ij} - \tilde{I}_{ij})|^2 + O(\|X - \tilde{I}\|_F^3)$$

$$(30) \quad \geq (\lambda_{m+1} - \lambda_m) \|X - \tilde{I}\|_F^2 + O(\|X - \tilde{I}\|_F^3).$$

The condition that $X_{ij} = X_{ji}$ was used to ensure that the $X_{ij} - \tilde{I}_{ij}$ and $X_{ji} - \tilde{I}_{ji}$ terms did not cancel each other out, and condition $\text{Im}(X_{jj}) = 0$ is necessary since the imaginary part of the diagonal entries does not occur in (29). Since the 3rd order terms are bounded near \tilde{I} , there exists a $\delta > 0$ such that for $d(X, \tilde{I}) < \delta$,

$$(31) \quad E_0(X) - \min E_0 \geq \alpha(\lambda_{m+1} - \lambda_m)d(X, \tilde{I})^2.$$

Since we assumed *without loss of generality* that $\tilde{I} \in \arg \min_{G \in \arg \min E_0} \|X - G\|_F$, we have actually proven the stronger statement: If $d(X, S_0) < \delta$, then

$$(32) \quad E_0(X) - \min E_0 \geq \alpha(\lambda_{m+1} - \lambda_m)d(X, S_0)^2.$$

By Lemma 2.5, there exists $\mu_0 > 0$ such that $0 \leq \mu < \mu_0$ implies $d(S_\mu, S_0) < \delta$. For such μ , (32) is satisfied with $X = X_\mu$ for any $X_\mu \in S_\mu$.

$$(33) \quad E_0(X_\mu) - \min E_0 \geq \alpha(\lambda_{m+1} - \lambda_m)d(X_\mu, S_0)^2.$$

We will now derive one final inequality which will complete the proof. First we observe that

$$(34) \quad E_0(X_\mu) + \mu \|X_\mu\|_1 = E_\mu(X_\mu) \leq E_0(X_0) = E_0(X_0) + \mu \|X_0\|_1,$$

for any $X_\mu \in S_\mu$ and $X_0 \in S_0$. If in particular, we choose $X_0 \in \arg \min_{G \in \arg \min E_0} \|X_\mu - G\|_F$, then we can show

$$(35) \quad \begin{aligned} E_0(X_\mu) - \min E_0 &\leq \mu (\|X_0\|_1 - \|X_\mu\|_1) \\ &\leq \mu \|X_0 - X_\mu\|_1 \\ &\leq \mu \|X_0 - X_\mu\|_0 \|X_0 - X_\mu\|_F \\ &\leq \mu \|X_0 - X_\mu\|_0 d(X_\mu, S_0), \end{aligned}$$

where $\|Y\|_0$ denotes the number of nonzero entries in Y . In particular, we can minimize $\|X_0 - X_\mu\|_0$ to obtain

$$(36) \quad E_0(X_\mu) - \min E_0 \leq \mu d(X_\mu, S_0) \cdot \min \left\{ \|X_0 - X_\mu\|_0 : X_0 \in \arg \min_{G \in \arg \min E_0} \|X_\mu - G\|_F \right\}.$$

We complete the proof by combining (33) with (36) and taking the max over $X_\mu \in S_\mu$.

$$(37) \quad d(S_\mu, S_0) \leq \frac{\mu}{\alpha(\lambda_{m+1} - \lambda_m)} \cdot \max_{X_\mu \in S_\mu} \min \left\{ \|X_0 - X_\mu\|_0 : X_0 \in \arg \min_{G \in \arg \min E_0} \|X_\mu - G\|_F \right\}$$

$$(38) \quad \leq \frac{\mu \sqrt{Nm}}{\alpha(\lambda_{m+1} - \lambda_m)}.$$

The proof is hence completed. We note that while (38) is a simple expression, there may exist problems for which (37) is also a tractable expression and provides a tighter bound. \square

We note that if $X_0 \in S_0$, then $\|X_0\| = \sqrt{m}$. So, the relative error does not depend explicitly on m . Theorem 3 also enables us to prove several corollaries. As a first application of the theorem, let us prove the convergence rate of $\min E_\mu$ and $E_0(X_\mu)$ to $\min E_0$ (where $X_\mu \in S_\mu$).

Corollary 1. *Let $\alpha \in (0, 1)$. For all $\mu > 0$, we have*

$$(39) \quad \min E_\mu - \min E_0 \leq \mu \min_{X_0 \in S_0} \|X_0\|_1.$$

And if $\lambda_m < \lambda_{m+1}$, then there is $\mu_0 > 0$ such that for $0 < \mu < \mu_0$,

$$(40) \quad E_0(X_\mu) - \min E_0 \leq \frac{\mu^2}{\alpha(\lambda_{m+1} - \lambda_m)} \cdot \min \left\{ \|X_0 - X_\mu\|_0^2 : X_0 \in \underset{G \in \arg \min E_0}{\arg \min} \|X_\mu - G\|_F \right\},$$

$$(41) \quad \leq \frac{\mu^2 Nm}{\alpha(\lambda_{m+1} - \lambda_m)}.$$

Proof. The first inequality is straightforward by noticing that if $X_\mu \in S_\mu$ and $X_0 \in S_0$, then $E_\mu(X_\mu) \leq E_\mu(X_0)$. The second inequality is obtained by combining (33) and (36). \square

We will perform numerical tests in Section 4.1 which suggest that the powers on μ in these bounds are optimal.

As a second corollary, we examine how orthonormal the columns of minimizers of E_μ are. The minimizers of E_0 are known to have orthonormal columns (Theorem 1), but the addition of the ℓ^1 term in E_μ discards this property. Nonetheless, we can still provide a bound on how close to orthonormal the columns are, as in the next corollary.

Corollary 2. *Let $\alpha \in (0, 1)$. Then, there exists $\mu_0(\alpha) > 0$ such that for $0 < \mu < \mu_0(\alpha)$ and $X_\mu \in S_\mu$, we have*

$$(42) \quad \left\| X_\mu^* X_\mu - I \right\|_F \leq \frac{2\mu m \sqrt{N}}{\alpha^2(\lambda_{m+1} - \lambda_m)}.$$

Proof. Choose μ small enough so that Theorem 3 applies and $d(S_\mu, S_0) \leq 2\sqrt{m}(\frac{1}{\alpha} - 1)$. Let $X_\mu \in S_\mu$. Let $X_0 \in \underset{X \in S_0}{\arg \min} \|X_\mu - X\|_F$. Then we can estimate,

$$(43) \quad \begin{aligned} \left\| X_\mu^* X_\mu - I \right\|_F &= \left\| X_\mu^* X_\mu - X_0^* X_0 \right\|_F \\ &\leq \left\| X_\mu^* (X_\mu - X_0) \right\|_F + \left\| (X_\mu^* - X_0^*) X_0 \right\|_F \\ &\leq (\|X_\mu\|_F + \|X_0\|_F) \|X_\mu - X_0\|_F \\ &\leq \frac{2\mu m \sqrt{N}}{\alpha^2(\lambda_{m+1} - \lambda_m)}, \end{aligned}$$

where the last line uses Theorem 3 and the fact that $\|X_0\|_F = \sqrt{m}$ since the columns of X_0 are orthonormal. \square

As our final corollary, we examine the convergence rate of the density matrix. Note that if we choose any $X_0 \in S_0$, then the density matrix $P_0 = X_0 X_0^*$ will be the projection operator onto the low-lying eigenspace. In what follows, if $X_\mu \in S_\mu$, then we will consider two different approximations for the density matrix.

$$(44) \quad \tilde{P}_{X_\mu} = X_\mu X_\mu^*,$$

$$(45) \quad P_{X_\mu} = X_\mu (X_\mu^* X_\mu)^{-1} X_\mu^*.$$

P_{X_μ} is a legitimate projection matrix while \tilde{P}_{X_μ} is approximately a projection matrix. The range of each is $\text{Col}(X_\mu)$. We can obtain the following estimates for the convergence of \tilde{P}_{X_μ} and P_{X_μ} to P_0 .

Corollary 3. *Let $\alpha \in (0, 1)$. Then, there exists $\mu_0(\alpha) > 0$ such that for $0 < \mu < \mu_0(\alpha)$, we have*

$$(46) \quad \left\| \tilde{P}_{X_\mu} - P_0 \right\|_F \leq \frac{2\mu m \sqrt{N}}{\alpha^2(\lambda_{m+1} - \lambda_m)},$$

$$(47) \quad \left\| P_{X_\mu} - P_0 \right\|_F \leq \frac{4\mu m \sqrt{N}}{\alpha^2(\lambda_{m+1} - \lambda_m)}.$$

Proof. Choose μ small enough so that Theorem 3 applies and $d(S_\mu, S_0) \leq 2\sqrt{m}(\frac{1}{\alpha} - 1)$. Let $X_\mu \in S_\mu$ and $X_0 \in \arg \min_{X \in S_0} \|X_\mu - X\|_F$. Then we can estimate,

$$(48) \quad \begin{aligned} \left\| \tilde{P}_{X_\mu} - P_0 \right\|_F &= \left\| X_\mu X_\mu^* - X_0 X_0^* \right\|_F \\ &\leq \left\| X_\mu (X_\mu^* - X_0^*) \right\|_F + \left\| (X_\mu - X_0) X_0^* \right\|_F \\ &\leq (\|X_\mu\|_F + \|X_0\|_F) \|X_\mu - X_0\|_F \\ &\leq \frac{2\mu m \sqrt{N}}{\alpha^2(\lambda_{m+1} - \lambda_m)}. \end{aligned}$$

A slightly more complicated calculation can be used to obtain (47).

$$(49) \quad \begin{aligned} \left\| P_{X_\mu} - P_0 \right\|_F &= \left\| X_\mu (X_\mu^* X_\mu)^{-1} X_\mu^* - X_0 X_0^* \right\|_F \\ &\leq \left\| X_\mu (X_\mu^* X_\mu)^{-1} X_\mu^* - X_\mu X_\mu^* \right\|_F + \left\| \tilde{P}_{X_\mu} - P_0 \right\|_F \\ &= \left\| I - X_\mu^* X_\mu \right\|_F + \left\| \tilde{P}_{X_\mu} - P_0 \right\|_F \\ &\leq \frac{4\mu m \sqrt{N}}{\alpha^2(\lambda_{m+1} - \lambda_m)}. \end{aligned}$$

The last line uses Corollary 2 and (46). We went from the second to the third line by use of

$$(50) \quad \begin{aligned} \left\| X_\mu (X_\mu^* X_\mu)^{-1} X_\mu^* - X_\mu X_\mu^* \right\|_F^2 &= \text{tr} \left[X_\mu \left[(X_\mu^* X_\mu)^{-1} - I \right] X_\mu^* X_\mu \left[(X_\mu^* X_\mu)^{-1} - I \right] X_\mu^* \right] \\ &= \text{tr} \left[X_\mu^* X_\mu \left[(X_\mu^* X_\mu)^{-1} - I \right] X_\mu^* X_\mu \left[(X_\mu^* X_\mu)^{-1} - I \right] \right] \\ &= \text{tr} \left[(I - X_\mu^* X_\mu)(I - X_\mu^* X_\mu) \right] \\ &= \left\| I - X_\mu^* X_\mu \right\|_F^2. \end{aligned}$$

□

2.3. Local minima of E_μ . In Theorem 2, we showed that all local minima of E_0 are global minima. Unfortunately, this is no longer true for E_μ . However, we can provide some information about the local minima of E_μ . The next theorem states that local minima of E_μ are in a sense “generated by” critical points of E_0 . Additionally, it states that these local minima of E_μ can be forced to be arbitrarily close to the critical points of E_0 by choosing μ sufficiently small. We denote the set of local minima of E_μ by S_μ^{loc} and denote the set of critical points of E_0 by C_p .

Theorem 4. *For every $\varepsilon > 0$, there exists $\delta > 0$ such that $0 \leq \mu < \delta$ implies $d(S_\mu^{\text{loc}}, C_p) < \varepsilon$.*

Proof. First we need a preliminary result: Let $X_{\min} \in S_{\mu}^{\text{loc}}$ for some $\mu > 0$. Let $\{x_i\}$ be the coordinate directions corresponding to the nonzero entries of X_{\min} and let $\{z_j\}$ be coordinate directions corresponding to the zero entries of X_{\min} . Then we must have

$$\begin{aligned} \frac{\partial E_0}{\partial x_i}(X_{\min}) &\in \{\pm\mu\} \\ \left| \frac{\partial E_0}{\partial z_j}(X_{\min}) \right| &\leq \mu \end{aligned}$$

for all i, j . This implies $\|\nabla E_0(X_{\min})\|_F \leq \mu\sqrt{Nm}$.

Next, let $U \in \mathbb{C}^{N \times m}$ such that $\|U\|_F = 1$. E_0 along the direction of U is given by

$$\begin{aligned} E_0(tU) &= -\text{tr}[U^* U U^* H U] t^4 + 2 \text{tr}[U^* H U] t^2 \\ (51) \quad &= c_1(U) t^4 + c_2(U) t^2. \end{aligned}$$

Since H is negative definite, $c_2(U) < 0 < c_1(U)$. So, on any line through the origin, E_0 has a double well potential shape. All critical points of E_0 must either be at the origin or at one of the minima along some line through the origin. By taking the derivative with respect to t , we see that if $Y \in C_p$, then $\|Y\|_F = 0$ or $\|Y\|_F = \sqrt{\frac{-c_2(Y/\|Y\|_F)}{c_1(Y/\|Y\|_F)}}$. Note that this expression is continuous with respect to Y for $Y \neq 0$. So, we can apply a compactness argument on the projective space to see that the set C_p is bounded.

Now to prove the lemma, we suppose by way of contradiction that the lemma is false. Then, there exists $\varepsilon > 0$ and sequences $\{\mu_k\}$ and $\{X_k\}$ such that $\mu_k \rightarrow 0$ and $X_k \in S_{\mu_k}^{\text{loc}}$ and $d(X_k, C_p) \geq \varepsilon$. First note that any local minima Y of E_{μ} must have a Frobenious norm less than $\sqrt{\frac{-c_2(Y/\|Y\|_F)}{c_1(Y/\|Y\|_F)}}$. This is clear because

$$(52) \quad E_{\mu}(tU) = c_1(U) t^4 + c_2(U) t^2 + \|U\|_1 \mu |t|.$$

Therefore, the set $\{X_k\}$ is bounded. So now by compactness, there exists a convergent subsequence $X_{k_l} \rightarrow Z$. Also, note that $\|\nabla E_0(X_k)\|_F \leq \mu_k \sqrt{Nm}$. Since $\mu_k \rightarrow 0$ and the gradient of E_0 is continuous, we have $\|\nabla E_0(Z)\|_F = 0$. That is, $Z \in C_p$. Therefore, there exists L such that $d(X_{k_L}, C_p) < \varepsilon$ which contradicts our assumption. \square

It is also worth commenting on the existence of local minima of E_{μ} . We note that the origin is a local minima of E_{μ} for all $\mu > 0$ (as can be seen from (52)). We also note that local minima can be generated for all $\mu > 0$ at points other than the origin. As an example, if $H = \begin{bmatrix} -1 & 0 \\ 0 & -2 \end{bmatrix}$ and $X \in \mathbb{C}^{2 \times 1}$, then the critical points $[e^{\alpha i}, 0]^T$ for all $\alpha \in \mathbb{R}$ generate a local minima of E_{μ} for all $\mu > 0$. However, for small μ , the basins around the local minima are small, so algorithms are less likely to converge to those points.

3. ALGORITHMS

Having determined many nice analytic properties of the functional E_{μ} , we now seek a practical way to find its minima using numerical optimization algorithms. We have two main objectives for our algorithms. First, we want them to find a minimum quickly. Second, we want them to maintain sparsity from iteration to iteration. The hope is that sparsity can be used to speed up the computational time per iteration through the use of sparse matrix arithmetic. We choose the Iterative Shrinkage–Thresholding Algorithm (ISTA) [1] as our base algorithm to

which we subsequently make modifications. This algorithm is a natural choice because it is designed to deal with objective functions which are the sum of a differentiable function (E_0) and a convex function (the ℓ^1 penalty term).

We will present a total of six different algorithms for minimizing E_μ . All of the algorithms are essentially modifications of ISTA. They can be divided into two main categories: block and non-block. The block algorithms update X column by column instead of updating all entries of X at once. Additionally, our algorithms could be categorized based on whether they use traditional backtracking or a more aggressive dynamic backtracking we propose. The numerical results in Section 4.2 show that the algorithms with this dynamic backtracking converge much faster than the algorithms using traditional backtracking.

3.1. Non-block version. ISTA with backtracking [1] is given in algorithm 1. In the algorithm, T_α is the shrinkage operator and \circ is an entrywise product. More specifically,

$$(53) \quad (T_\alpha Y)_{ij} = \begin{cases} 0 & \text{if } |Y_{ij}| \leq \alpha, \\ (|Y_{ij}| - \alpha) e^{i \arg(Y_{ij})} & \text{if } |Y_{ij}| > \alpha. \end{cases}$$

$$(54) \quad A \circ B = \sum_i \sum_j A_{ij} B_{ij}.$$

where $\arg(z)$ denotes the complex argument of z . The algorithm is essentially a steepest descent (corresponding to E_0) algorithm combined with an entrywise soft thresholding. The backtracking is used to choose the step size, which will be further explained below.

Algorithm 1 ISTA with backtracking

Input: $X_0 \in \mathbb{R}^{N \times m}$, $L_0 > 0$, $\eta > 1$

- 1: **for** $k = 1, 2, \dots$ **do** ▷ Loop until convergence is achieved
 - 2: $L_k^{(1)} \leftarrow L_{k-1}$
 - 3: **for** $j = 1, 2, \dots$ **do** ▷ Backtracking iteration
 - 4: $X_k^{(j)} \leftarrow T_{\mu/L_k^{(j)}} \left(X_{k-1} - \frac{1}{L_k^{(j)}} \nabla E_0(X_{k-1}) \right)$
 - 5: $L_k^{(j+1)} \leftarrow \eta L_k^{(j)}$
 - 6: **until** $E_0(X_k^{(j)}) \leq E_0(X_{k-1}) + \nabla E_0(X_{k-1}) \circ (X_k^{(j)} - X_{k-1}) + \frac{L_k^{(j)}}{2} \|X_k^{(j)} - X_{k-1}\|_F^2$
 - 7: $X_k \leftarrow X_k^{(j)}$
 - 8: $L_k \leftarrow L_k^{(j)}$
 - 9: **end for**
-

The condition $\|X_k - X_{k-1}\|_F < \text{tol}$ can be easily used as a convergence criteria as the quantity $\|X_k - X_{k-1}\|_F$ needs to be calculated anyway.

3.1.1. Dynamic backtracking. Numerical results indicated that the convergence could be quite slow using the backtracking version of ISTA. So we improved the algorithm by choosing L_k in a more dynamic fashion. At the start of each iteration, we choose the value of $L_k^{(1)}$ based on the previous iterations. Backtracking is still used if necessary.

To see the idea behind this modification, it is necessary to recall the derivation of ISTA. In particular, the basic idea is to replace the objective function E_μ with an approximation \tilde{E}_μ of

the objective function near X_{k-1} ,

$$(55) \quad \tilde{E}_\mu(X_{k-1}; X) = E_0(X_{k-1}) + \nabla E_0(X_{k-1}) \circ (X - X_{k-1}) + \frac{L_k}{2} \|X - X_{k-1}\|_F^2 + \mu \|X\|_1.$$

Then each step of ISTA is just given by $X_k = \operatorname{argmin} \tilde{E}_\mu(X_{k-1}; X)$. So, we can see that if L_k is a good approximation to the second directional derivative in the direction $X - X_{k-1}$, then \tilde{E}_μ would be a good approximation of E_μ in that direction. For this reason, we choose $L_k^{(1)}$ by approximating the second directional derivative. In particular, when $k \geq 2$, we make the following modification to line 2 of algorithm 1.

$$2: L_k^{(1)} \leftarrow c_1 \frac{\|\nabla E_0(X_{k-1}) - \nabla E_0(X_{k-2})\|_F}{\|X_{k-1} - X_{k-2}\|_F}$$

where $c_1 > 1$ (we found that $c_1 = 1.5$ worked well). The constant c_1 is used to prevent our guess for $L_k^{(1)}$ from being too small (which would lead to backtracking). This guess for $L_k^{(1)}$ will typically be much smaller than if we used the traditional backtracking approach $L_k^{(1)} = L_{k-1}$. This allows us to take larger step sizes which leads to faster convergence.

The drawbacks of this method are the extra computational cost associated with calculating $L_k^{(1)}$ and the extra cost of having to do more backtracking steps. However, our numerical results show that these extra costs are worthwhile because the accelerated convergence leads to many fewer iterations being required.

Furthermore, we will now provide another modification to the algorithm which greatly reduces the number of backtracks that are required. We can reduce the number of backtracks by modifying line 5 in the algorithm. If line 2 happens to give a guess for L_k which is way too small, then multiple backtracks will be required in a single iteration. To reduce the occurrence of multiple backtracks, we can choose $L_k^{(j+1)}$ in a smarter way. In particular, we can replace line 5 by

$$5: L_k^{(j+1)} \leftarrow c_2 \frac{2 \left[E_0(X_k^{(j)}) - E_0(X_{k-1}) - \nabla E_0(X_{k-1}) \circ (X_k^{(j)} - X_{k-1}) \right]}{\|X_k^{(j)} - X_{k-1}\|_F^2}$$

where $c_2 > 1$ (we used $c_2 = 2$). Notice that this calculation is essentially “free” since all these terms need to be calculated anyway for the break condition in line 6. The new line 5 is obtained by simply solving for $L_k^{(j)}$ in the break condition (and then multiplying by a constant c_2 greater than 1 to try to avoid multiple backtracks per iteration). In practice, using $c_2 > 1.1$ seemed to virtually eliminate the occurrence of multiple backtracks per iteration.

3.2. Block version. We now present our block algorithm. This algorithm (with dynamic backtracking) is given in algorithm 2. In order to implement it, we must first choose how to split up X into blocks. We choose each column of X to be its own block. This gives a total of m blocks. By choosing the blocks this way, m iterations of the block algorithm has essentially the same cost as one iteration of the non-block algorithm.

The following notation is used in algorithm 2. X_{ℓ, b_k} is the b_k block of X_ℓ , and ∇_{b_k} is the gradient with respect to the b_k block, holding all other blocks constant. We also make use of the function

$$(56) \quad p(b_k) = \begin{cases} \max\{\ell < k : b_\ell = b_k\} & \text{if } \{\ell < k : b_\ell = b_k\} \neq \emptyset, \\ 0 & \text{if } \{\ell < k : b_\ell = b_k\} = \emptyset. \end{cases}$$

Algorithm 2 Block version with dynamic backtracking

Input: $X_0 \in \mathbb{R}^{N \times m}$, $L_1^{(1)} > 0$, $c_1, c_2 > 1$

```

1: for  $k = 1, 2, \dots$  do                                     ▷ Loop until convergence is achieved
2:   Choose  $b_k \in \{1, 2, \dots, m\}$ .
3:    $L_k^{(1)} \leftarrow c_1 \frac{\|\nabla_{b_k} E_0(X_{k-1}) - \nabla_{b_k} E_0(X_{p(b_k)-1})\|_F}{\|X_{k-1, b_k} - X_{p(b_k)-1, b_k}\|_F}$    (if  $k \geq 2$ )
4:   for  $j = 1, 2, \dots$  do                                     ▷ Backtracking iteration
5:      $X_{k, b_k}^{(j)} \leftarrow T_{\mu/L_k^{(j)}} \left( X_{k-1, b_k} - \frac{1}{L_k^{(j)}} \nabla_{b_k} E_0(X_{k-1}) \right)$ 
6:      $L_k^{(j+1)} \leftarrow c_2 \frac{2 \left[ E_0(X_k^{(j)}) - E_0(X_{k-1}) - \nabla_{b_k} E_0(X_{k-1}) \circ (X_{k, b_k}^{(j)} - X_{k-1, b_k}) \right]}{\|X_{k, b_k}^{(j)} - X_{k-1, b_k}\|_F^2}$ 
7:   until  $E_0(X_k^{(j)}) \leq E_0(X_{k-1}) + \nabla_{b_k} E_0(X_{k-1}) \circ (X_{k, b_k}^{(j)} - X_{k-1, b_k}) + \frac{L_k^{(j)}}{2} \|X_{k, b_k}^{(j)} - X_{k-1, b_k}\|_F^2$ 
8:    $X_k \leftarrow X_{k-1}$ 
9:    $X_{k, b_k} \leftarrow X_{k, b_k}^{(j)}$                                ▷ Only the  $b_k$  block changes
10: end for
    
```

That is, $p(b_k)$ is the most recent iteration that the b_k block was updated, and if the b_k block has not yet been updated, then we define $p(b_k) = 0$.

There are clearly many possible ways to choose b_k in line 2 of each iteration, but for our numerical tests in Section 4, we will only consider two such strategies. In our “sequential” strategy, we simply loop sequentially through the blocks $(1, 2, \dots, m, 1, 2, \dots)$. In what we call the “random” strategy for choosing b_k , we choose a random permutation of $\{1, \dots, m\}$ and then use that for the first m iterations. For the next m iterations, we choose a new random permutation of $\{1, \dots, m\}$, and so on. In this way, each block is updated exactly once between iterations $(q-1)m+1$ and qm for any $q \in \mathbb{N}$. With our choices of b_k , we note that m iterations of the block algorithm is approximately the same as one iteration of the non-block algorithm (since each block is updated exactly once). So, when we compare the two methods, we think of m iterations of the block algorithm as being a single iteration.

4. NUMERICAL RESULTS

In this section we perform numerical tests to verify our theoretical results (Section 4.1), compare the performance of our proposed algorithms (Section 4.2), and explore some of the more practical concerns of our method (Sections 4.3, 4.4 and 4.5). We will use the following example, borrowed from [4], as our test problem in the remainder of the paper.

$$(57) \quad H = -\frac{1}{2} \frac{d^2}{dx^2} + V(x) - \eta,$$

$$(58) \quad V(x) = \alpha \sum_j e^{-(x-r_j)^2/(2\beta^2)},$$

where η is a constant chosen so that H is negative definite. The domain of the problem is $[0, 10]$ with periodic boundary conditions, and $\{r_j\}_{j=1}^{10} = \{0.5, 1.5, \dots, 9.5\}$ represent the atom positions. We discretize on a uniform spacial grid and use the centered difference formula to discretize

μ	$\min E_\mu - \min E_0$	order	$E_0(X_\mu) - \min E_0$	order	$d(X_\mu, S_0)$	order
2^{-8}	2.4412e-01	–	7.5520e-05	–	1.1147e-03	–
2^{-9}	1.2208e-01	0.99976	2.1369e-05	1.8213	5.9890e-04	0.89628
2^{-10}	6.1045e-02	0.99987	5.8575e-06	1.8672	3.1342e-04	0.93420
2^{-11}	3.0524e-02	0.99993	1.6183e-06	1.8558	1.6449e-04	0.93012
2^{-12}	1.5262e-02	0.99996	4.4340e-07	1.8678	8.5349e-05	0.94653

TABLE 1. Large gap problem: Convergence with respect to μ .

μ	$\min E_\mu - \min E_0$	order	$E_0(X_\mu) - \min E_0$	order	$d(X_\mu, S_0)$	order
2^{-8}	4.5203e-01	–	2.6667e-03	–	2.0772e-02	–
2^{-9}	2.2668e-01	0.99576	6.9602e-04	1.9853	1.0505e-02	0.9835
2^{-10}	1.1351e-01	0.99783	1.9228e-04	1.9431	5.4162e-03	0.9556
2^{-11}	5.6799e-02	0.99891	5.3083e-05	1.8623	2.7386e-03	0.9838
2^{-12}	2.8410e-02	0.99946	1.3839e-05	1.9264	1.2868e-03	1.0896

TABLE 2. Small gap problem: Convergence with respect to μ .

the second derivative. We take $m = 10$, and the initial condition X_0 is set as follows. The i^{th} column of X_0 has support of width $2L + 1$ centered at r_i . On the support, the initial condition is set by generating random numbers between 0 and $2/(2L + 1)$ so that the expected 2–norm of each column is 1. Unless otherwise specified, we will use the parameters $\alpha = -100$ and $\beta = 0.1$. In this case, the spectral gap is $\lambda_{m+1} - \lambda_m \approx 54.2$.

4.1. Convergence with respect to μ . The first test we perform is to confirm the theoretical results in Theorem 3 and Corollary 1. We list the minimal value of E_μ (and the value of E_0 at the minimum) compared to the minimal value of E_0 to validate Corollary 1. We also investigate the distance from a minimizer of E_μ to the set S_0 to test the linear convergence rate guaranteed by Theorem 3.

To perform this test, we take $N = 800$, and use two different choices for α which give us different spectral gaps. In particular, we choose $\alpha = -100$ and $\alpha = -10$ which give gaps of 54.2 and 4.36 respectively.

The results for different values of μ are tabulated in Tables 1 and 2. From these results, we see that the powers of μ in our theoretical results appear to be optimal. We also see that the problem with a smaller gap converges slower than the one with a larger gap in accordance with the constant out front in Theorem 3 and Corollary 1.

To get an intuitive view of what is happening to the solution as we change μ , we have plotted the 5th column of X in Figure 1. From this figure, we can see the localizing behavior of the orbital as μ gets larger. We also see that for small μ , the orbital is not much different than when $\mu = 0$.

4.2. Comparing convergence of algorithms. Here we test the convergence rates of our proposed algorithms. The results in Figure 2 are typical when the initial conditions have small support, *i.e.*, are better approximations of the solution, since the solution decays away from the r_i 's. The dependence on initial conditions will be discussed further in Section 4.4. We can see that the dynamic backtracking algorithms converge much faster than the traditional backtracking algorithms. This increase in speed is due to the fact that the dynamic algorithms can




FIGURE 1. Here the 5th column of X is plotted for different values of μ . The result from the large gap problem is zoomed in for a better view.




FIGURE 2. Convergence rate test

use much larger step sizes than the traditional algorithms. The convergence rate appears to be linear, and there does not appear to be much of a difference between the required number of iterations for the block and non-block versions.

4.3. Local minima test. Our method is not the first to try and take advantage of the “almost compact” structure of the orbitals. A previous attempt was the truncated steepest descent (truncated SD) method. However, it is known to have the unfortunate property of converging to local minima [4]. Since our functional E_μ is also known to contain local minima, we wish to test how often our proposed method gets trapped in local minima compared to the truncated SD method. Our method has an additional parameter μ whose effect also needs to be tested. It is intuitive that as μ gets larger, it creates larger “basins” around the local minima of E_μ . So we expect our algorithms to get stuck at local minima more often when μ is larger.

The following test compares many things: truncated SD vs our proposed methods, block vs non-block, and the effect of using different values of μ . Our test setup is as follows. For a given initial condition, we run the truncated SD method as well as the block and non-block dynamic backtracking methods with different values for μ . We use the large gap problem with $N = 500$ and $L = 60$. All these methods have different global minima of their respective energy functionals, so to compare them all we plot (Figure 3) how far above this minimum the respective algorithms converged to. That is, for the $\mu = 10$ test, we plot $E_{10}(X) - \min E_{10}$. We do likewise for the other values of μ . For the truncated SD method, we plot $E_0(X) - \min E_0$.

We can see from these results that even though our method is known to sometimes have local minima of E_μ , it doesn't seem to be an issue in our test problem for small μ . Our method avoids local minima much better than truncated SD for small μ , but gets stuck in local minima much more often for large μ . We also note that neither the block nor the non-block version of the algorithm seems to be much better or worse at avoiding local minima. We also ran the same test with 10,000 trials for the non-block version with $\mu = 0.5$, and the algorithm did not get stuck at a local minima even once.

In our results, we notice that when the algorithm gets stuck at local minima, these local minima are essentially grouped into discrete energy levels. This is especially seen well in the non-block version. This behavior is expected because the local minima of E_μ are known to occur near critical points of E_0 (Theorem 4) which are spans of eigenvectors of H (Lemma 2.3). Therefore, we expect that the energy value at local minima will be approximately the sum of some m or fewer eigenvalues of H . This is what leads to the observed discrete energy levels in the plots.




FIGURE 3. Plots of the energy (above the minimum) that the algorithms converged to. Each column corresponds to a distinct initial condition and the initial conditions are the same in both plots.

4.4. Dependence on initial conditions. In our next numerical test we investigate how the number of iterations required to converge depends on the initial condition. In general, it would be expected that a “good” initial guess would lead to a smaller number of iterations required for convergence and a “bad” initial guess would lead to a larger number of iterations. We will see that this is true, but there is also more to the story. While we want to converge in as few iterations as possible, we also want to preserve sparsity from iteration to iteration. By having a sparse X , we hope that the calculations can be done faster using sparse matrix multiplication

algorithms. We will see that this idea of preserving sparsity is actually intimately related to the number of iterations required to converge.

In order to get a better idea of what is going on, we note that sparsity from iteration to iteration is essentially determined in the following step of the algorithm.

$$X_k^{(j)} \leftarrow T_{\mu/L_k^{(j)}} \left(X_{k-1} - \frac{1}{L_k^{(j)}} \nabla E_0(X_{k-1}) \right)$$

The part in the parentheses tends to increase the number of nonzero entries while the shrinkage operator tends to decrease them. It is clear just from this formula that a larger μ will tend to preserve sparsity better than a smaller μ . However, there are two critical reasons for wishing to avoid using a large μ . The first is that the minima of E_μ will be farther away from the minima of E_0 . The second (as seen in Section 4.3) is that the basins around the local minima of E_μ will be larger which increases the chance of converging to a local minima of E_μ .

We will use the dynamic backtracking version of ISTA for this test. In Figure 4, we plot the number of iterations that each entry of X was nonzero. In particular, our test problem uses the parameters $N = 150$, $m = 10$, and $\mu = 0.1$. The only difference between the tests is the support of the initial conditions. The support of the initial condition acts as a surrogate for how “good” the initial condition is. With the way we are generating the initial conditions, a smaller support would tend to be closer in shape to the actual solution, and therefore a “better” initial guess. We see from Figure 5 that the smaller the value of L , the fewer iterations that are required for convergence.

For $L > 4$, we can see plateaus in the plots of Figure 5. It is worthwhile to explore what is going on here since these plateaus dramatically increase the number of iterations required to converge. We notice that for larger L , the algorithm actually converges to approximately the minimal value of E_0 very quickly, but then struggles to decrease its ℓ^1 norm. During the plateau region, the algorithm is slowly decreasing its ℓ^1 norm while keeping its E_0 value fairly constant. In other words, by the beginning of the plateau, the algorithm has found a point very near S_0 , but this point has a large ℓ^1 norm. So, the plateau is essentially spent rotating (recall that E_0 is invariant under rotations) the solution into regions of lower ℓ^1 norm that are also near S_0 .

To be more concrete, in our tests we have found that these plateau regions are caused by overlap of the orbitals (the ℓ^1 norm will be smaller when they are non-overlapping). We can see this in Figure 4. For example, in the $L = 12$ case, orbitals 6 and 7 overlapped for much of the simulation. It is as if a Givens rotation had been applied to the 6th and 7th orbitals of an exact minimizer of E_μ . When two (or more) orbitals are “competing” with each other like this, convergence tends to be very slow until the competition is resolved (ie. one of them wins by essentially zeroing out the other one). This is what causes those long plateaus in the graph of error vs iterations when $L = 8, 12, 16$. So, it is important to avoid overlap because it leads to both less sparsity (more time per iteration) and slower convergence (more iterations).

4.5. Dynamic μ . The above discussions on local minima, Theorem 3, and maintaining sparsity lead us to some competing objectives for our choice of μ . On the one hand, we want a small μ so that the minima of E_μ are close to the minima of E_0 and so that we can avoid being trapped at local minima. On the other hand, we want a large enough μ so that we can maintain sparsity which leads to faster convergence. In this section, our goal is to show the possible benefit of changing μ from iteration to iteration. We do not present a robust algorithm, but rather just a numerical test which provides hope that future work in this area could be beneficial.




FIGURE 4. The above plots show the number of iterations for which each particular entry of X was nonzero. The layout of the plots corresponds to the transpose of X . Note the different scales for each plot.

In general, the main idea for choosing μ at each iteration should be based on the following points. First, μ should be small enough at the beginning of the simulation so that the algorithm has a very low probability of getting stuck at a local minimum. Second, after the algorithm has found a near minimal value of E_0 , the value of μ could be increased in order to facilitate reduction of the ℓ^1 norm. A large μ provides a stronger “force” pushing X towards a lower ℓ^1 norm. Note that using a larger μ will also push X towards larger values of E_0 . However, as long as μ is not increased too much, we expect the value of E_0 to remain below $\min E_0 + \lambda_{m+1} - \lambda_m$. Staying below this value would be preferred as we know from Theorem 4 and Lemma 2.3 that for small μ , local minima are near values of E_0 that correspond to the sum of some m eigenvalues of H . So, we would hope not to get stuck at a local minimum if we only moderately increase μ once we are already near a minimal value of E_0 . Finally, if we want our solution to be close to




FIGURE 5. Convergence of the dynamic backtracking version of ISTA for different L . The plots here correspond to the sparsity plots in Figure 4. Solid lines plot $E_\mu(X) - \min E_\mu$. Dashed lines plot $E_0(X) - \min E_0$. Note that we do not expect the error in E_0 to go to 0 since that would only happen in the case $\mu = 0$.

the exact solution, we want the value of μ to be small. So, after the ℓ^1 norm has been sufficiently decreased, μ could also be decreased to allow for a solution closer to an exact minimizer of E_0 .

Of course, such a strategy is not so straightforward as one does not usually know a priori the values of $\min E_0$ and $\min E_\mu$. So, we leave it as an open problem how one could implement such ideas in practice. Nevertheless, we perform a numerical test which justifies the idea of increasing μ after several iterations. To illustrate the effect that changing μ in this manner can have, we will carry out the above problem with $L = 16$ again. We run the algorithm twice (starting with the same initial condition). One run will use a constant μ value of 0.1. The other will also use a μ value of 0.1 except for iterations 100–499 during which μ will be set to 1. We can see from Figure 6 that the algorithm converges much faster when we change μ in this way (≈ 550 iterations instead of ≈ 2900). It is an open problem how this type of strategy could be carried out in a more dynamic and intelligent manner.

5. DISCUSSION

We have presented a new method for OMM calculations which takes advantage of almost compact orbitals by introducing an ℓ^1 penalty term into the energy functional. Our analysis proves the convergence of the minimizers of E_μ to the minimizers of E_0 which justifies the practice of minimizing E_μ rather than E_0 . Numerical results show the ability of our algorithms to maintain sparsity from iteration to iteration. This suggests that our algorithms combined with sparse matrix algorithms could have high performance. Implementing such a scheme and comparing with other schemes which minimize the OMM energy functional is a direction for future work.

There are several additional possibilities for future work. One is to prove convergence results for our proposed algorithms. Analysis has been done on algorithms similar to the ones we present with traditional backtracking [22, 19]. It will be interesting to investigate the convergence of the dynamic backtracking algorithms. Another possible future direction would be to parallelize the block algorithm: all the blocks would be updated at the same time (in parallel)




FIGURE 6. Convergence with constant μ versus a variable μ . Note the different scales in the two plots. Also note that for $\mu = 0.1$, $E_0(X_\mu) \approx \min E_0 + 5.9 \cdot 10^{-3}$. In plot 6b, we see the value of $E_0(X)$ jump up when μ increases and jump down when μ decreases. This is not surprising since a larger μ would take the solution farther away from the minimum value of E_0 .

rather than one after another. The convergence of such an algorithm needs further investigation. A further direction would be to explore the use of a dynamic μ , as already discussed in Section 4.5. Last but not least, the convergence of the algorithm might be accelerated by using a well chosen preconditioner, for example as considered in the recent work [11].

REFERENCES

- [1] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [2] F. Corsetti. The orbital minimization method for electronic structure calculations with finite-range atomic basis sets. *Comput. Phys. Commun.*, 185:873–883, 2014.
- [3] W. E, T. Li, and J. Lu. Localized bases of eigensubspaces and operator compression. *Proc Natl Acad Sci USA*, 107(1273–1278), 2010.
- [4] W. Gao and W. E. Orbital minimization with localization. *Discrete Contin. Dyn. Sys.*, 23:249–264, 2009.
- [5] S. Geodecker. Linear scaling electronic structure methods. *Rev. Mod. Phys.*, 71:1085–1123, 1999.
- [6] J. Kim, F. Mauri, and G. Galli. Total-energy global optimizations using nonorthogonal localized orbitals. *Physical Review B*, 52(3):1640, 1995.
- [7] W. Kohn. Analytic Properties of Bloch Waves and Wannier Functions. *Physical Review*, 115(4):809–821, 1959.
- [8] R. Lai and J. Lu. Localized density matrix minimization and linear scaling algorithms. *arXiv preprint arXiv:1506.01610*, 2015.
- [9] R. Lai, J. Lu, and S. Osher. Density matrix minimization with ℓ_1 regularization. *Commun. Math. Sci.*, 13(5):1051–1074, 2015.
- [10] R. Lai and S. Osher. A splitting method for orthogonality constrained problems. *Journal of Scientific Computing*, 58(2):431–449, 2014.

- [11] J. Lu and H. Yang. Preconditioning orbital minimization method for planewave discretization, 2016. preprint.
- [12] N. Marzari and D. Vanderbilt. Maximally localized generalized Wannier functions for composite energy bands. *Physical Review B*, 56(20):12847–12865, 1997.
- [13] F. Mauri and G. Galli. Electronic-structure calculations and molecular-dynamics simulations with linear system-size scaling. *Physical Review B*, 50(7):4316, 1994.
- [14] F. Mauri, G. Galli, and R. Car. Orbital formulation for electronic-structure calculations with linear system-size scaling. *Physical Review B*, 47(15):9973, 1993.
- [15] P. Ordejón, D. A. Drabold, M. P. Grumbach, and R. M. Martin. Unconstrained minimization approach for electronic computations that scales linearly with system size. *Phys. Rev. B*, 48:14646–14649, 1993.
- [16] P. Ordejón, D. A. Drabold, R. M. Martin, and M. P. Grumbach. Linear system-size scaling methods for electronic-structure calculations. *Phys. Rev. B*, 51:1456–1476, 1995.
- [17] V. Ozolins, R. Lai, R. Caflisch, and S. Osher. Compressed modes for variational problems in mathematics and physics. *Proc. Natl. Acad. Sci. USA*, 110:18368–18373, 2013.
- [18] B. G. Pfrommer, J. Demmel, and H. Simon. Unconstrained energy functionals for electronic structure calculations. *Journal of Computational Physics*, 150(1):287–298, 1999.
- [19] P. Tseng and S. Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming*, 117(1-2):387–423, 2009.
- [20] E. Tsuchida. Augmented orbital minimization method for linear scaling electronic structure calculations. *Journal of the Physical Society of Japan*, 76:034708, 2007.
- [21] G. H. Wannier. The structure of electronic excitation levels in insulating crystals. *Physical Review*, 52(3):0191–0197, Aug. 1937.
- [22] Y. Xu and W. Yin. A globally convergent algorithm for nonconvex optimization based on block coordinate update. *arXiv preprint arXiv:1410.1386*, 2014.

DEPARTMENT OF MATHEMATICS, DEPARTMENT OF PHYSICS, AND DEPARTMENT OF CHEMISTRY, DUKE UNIVERSITY, BOX 90320, DURHAM NC 27708, USA

E-mail address: jianfeng@math.duke.edu

DEPARTMENT OF MATHEMATICS, DUKE UNIVERSITY, BOX 90320, DURHAM NC 27708, USA

E-mail address: kyle.thicke@duke.edu