

Score-based diffusion priors for multi-target detection

Alon Zabatani

School of Electrical Engineering
Tel Aviv University
alonzabatani@mail.tau.ac.il

Shay Kreymer

School of Electrical Engineering
Tel Aviv University
shaykreymer@mail.tau.ac.il

Tamir Bendory

School of Electrical Engineering
Tel Aviv University
bendory@tauex.tau.ac.il

Abstract—Multi-target detection (MTD) is the problem of estimating an image from a large, noisy measurement that contains randomly translated and rotated copies of the image. Motivated by the single-particle cryo-electron microscopy technology, we design data-driven diffusion priors for the MTD problem, derived from score-based stochastic differential equations models. We then integrate the prior into the approximate expectation-maximization algorithm. In particular, our method alternates between an expectation step that approximates the expected log-likelihood and a maximization step that balances the approximated log-likelihood with the learned log-prior. We show on two datasets that adding the data-driven prior substantially reduces the estimation error, in particular in high noise regimes.

Index Terms—Diffusion models, multi-target detection, cryo-EM, score-SDE, expectation-maximization

I. INTRODUCTION

We study the two-dimensional multi-target detection (MTD) problem: estimating an image $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ from a single, noisy measurement that contains multiple copies of the image, each randomly rotated and translated [1], [2]. In particular, we model the measurement $M \in \mathbb{R}^{N \times N}$ as

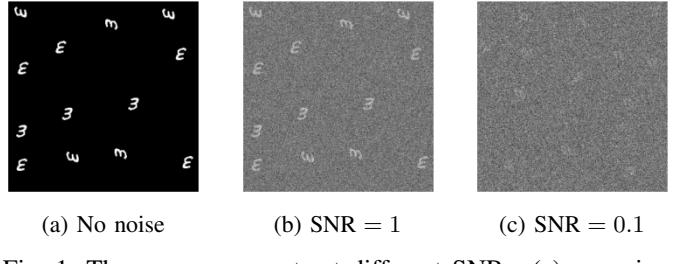
$$M[\vec{\ell}] = \sum_{i=1}^p F_{\phi_i}[\vec{\ell} - \vec{\ell}_i] + \varepsilon[\vec{\ell}], \quad (1)$$

where $F_{\phi_i}[\vec{\ell}] = f_{\phi_i}(\vec{\ell}/n)$ is a discrete copy of f , rotated by angle ϕ_i about the origin, n is the known radius of the image in pixels, $\{\phi_i\}_{i=1}^p$ and $\{\vec{\ell}_i\}_{i=1}^p$ are the unknown rotations and translations, respectively, and $\varepsilon[\vec{\ell}]$ is i.i.d. Gaussian noise with zero mean and variance σ^2 .

The goal is to estimate the image f from measurement M , while the rotations, translations, and the number of occurrences of f in M , denoted by p , are unknown. Importantly, since the rotations are unknown, it is possible to reconstruct the target image only up to a rotation. For simplicity, we assume that the translations are not densely packed, namely,

$$|\vec{\ell}_i - \vec{\ell}_j| > 4n, \quad \text{for all } i \neq j. \quad (2)$$

This assumption was also taken by previous works [1], [2] for simplification. The algorithm presented in this paper can be readily extended to handle arbitrarily spread signal occurrences, which violate (2), along with the guidelines that were proposed in [3], [4]. Figure 1 illustrates measurements with the digit three, taken from the MNIST dataset [5], at different



(a) No noise (b) SNR = 1 (c) SNR = 0.1

Fig. 1: Three measurements at different SNRs: (a) no noise; (b) SNR = 1; (c) SNR = 0.1. Each measurement contains multiple rotated versions of the digit three, taken from the MNIST dataset. We are primarily interested in low SNR regimes (e.g., panel (c)), in which it is extremely challenging to estimate the locations and rotations associated with the image occurrences.

signal-to-noise ratios (SNRs). We define $\text{SNR} \triangleq \frac{\|F\|_F^2}{A\sigma^2}$, where A is the area in pixels of F .

The MTD model is mainly studied as a mathematical abstraction of the single-particle cryo-electron microscopy (cryo-EM) technology for determining macromolecular structures [6]–[8]. In a cryo-EM experiment, individual copies of a target biomolecule, rotated by unknown 3-D rotations, are scattered across unknown 2-D locations within a thin layer of vitreous ice, from which 2-D tomographic projection images are produced by an electron microscope [9]. Maintaining a low electron dose is imperative to prevent irreversible structural damage, leading to significantly noisy projection images.

The conventional cryo-EM data processing pipeline [10]–[13] involves the initial detection and extraction of the 2-D projection images from the measurement, followed by their rotational and translational alignment to reconstruct the 3-D molecular structure. However, this approach fails for small molecules, which are too difficult to detect and align [6], [10], [14], [15]. Consequently, the MTD model—that can be thought of as a simplification of the cryo-EM model, where 2-D rotations replace the 3-D rotations, and the tomographic projection operator is neglected—was devised in [15], [16], to study the recovery of small molecular structures, without intermediate detection, below the current detection limit [17].

Previous works proposed an expectation-maximization (EM) algorithm to solve the MTD problem [3], [18], out-

performing an autocorrelation analysis alternative [1], [2], [4]. EM is an iterative algorithm to find maximum likelihood (or maximum a posteriori) estimators, where the underlying model depends on unobserved nuisance parameters (in our case, the rotations and translations associated with the p image occurrences). EM alternates between an expectation step (E-step), which calculates the expected log-likelihood evaluated using the current estimate of the parameters, and a maximization step (M-step), which finds the parameters maximizing the function computed in the E-step [19]. While EM cannot be applied directly to the MTD problem, due to a large number of possible translations that should be taken into account, the authors of [3] suggested an approximation scheme that dramatically reduces the computational burden, while providing high estimation quality; see Section III for details.

Motivated by recent attempts to utilize prior information to constitute molecular structures in cryo-EM (see for example [20], [21]), we suggest incorporating score-based diffusion priors into the approximate EM framework. Diffusion models [22]–[25] are a class of probabilistic generative models that progressively perturb data by injecting noise, then learn to reverse this process, thus capturing the implicit prior of the underlying data distribution. Diffusion models have been successfully applied to a variety of tasks [26]–[29], such as image synthesis, where they have broken the long-time dominance of generative adversarial networks (GANs) [30]. Score-based stochastic differential equations (score-SDE) [25] is a popular diffusion model, that makes use of the Stein score—the gradient of the log probability density [31], [32]. The key idea behind score-SDE is to perturb the data with a sequence of intensifying Gaussian noise governed by a closed-form stochastic differential equation, and jointly estimate the score functions for all noisy data distributions by training a deep neural network, conditioned on the noise levels.

The main contribution of this paper is the integration of score-based diffusion priors into the approximate EM framework, resulting in a flexible algorithm that balances between the log-likelihood term (the data) and the log-prior term using a regularized gradient ascent algorithm. We demonstrate on two datasets that we can successfully apply it to the two-dimensional MTD problem (1) and significantly improve image recovery quality, including in highly noisy regimes.

The rest of the paper is organized as follows. Section II introduces the score-SDE method. Section III presents the integration of score-SDE priors into the approximate EM framework. Section IV provides numerical results on two datasets (Gaussian images and MNIST), and Section V concludes the paper and delineates future applications to cryo-EM.

II. SCORE-BASED DIFFUSION MODELS

Score-based diffusion models learn the implicit prior of the underlying data distribution by matching the gradient of the log probability density with respect to data point x , i.e., Stein score function [31], [32]

$$s(x) \triangleq \nabla_x \log p(x). \quad (3)$$

We approximate the score function using a neural network $s_\theta(x)$, where the network's parameters θ are being learned by minimizing the following score matching loss:

$$\text{loss}(\theta) = \mathbb{E}_{x \sim p(x)} \|s_\theta(x) - \nabla_x \log p(x)\|_2^2. \quad (4)$$

However, in practice, the true gradients of the log probability density functions are often unknown during training. In [32], it was shown that the loss function (4) is equivalent to

$$\text{loss}(\theta) = \mathbb{E}_{x \sim p(x)} \left(\frac{1}{2} \|s_\theta(x)\|_2^2 + \text{tr}(\nabla_x s_\theta(x)) \right), \quad (5)$$

where $\text{tr}(\cdot)$ is the trace of the Jacobian matrix $\nabla_x s_\theta(x)$. Unfortunately, learning the score function by minimizing the loss function (5) tends to fail in regions where there are fewer training samples (due to the expectation over the data density); this can be mitigated by perturbing the data with various noise levels, as shown in [24]. Specifically, score-SDE [25] perturbs the data with a sequence of intensifying Gaussian noise, governed by a closed-form stochastic differential equation, and jointly estimates the score functions for all noisy data distributions. The SDE is defined as follows

$$dx = -\frac{\beta(t)}{2} x dt + \sqrt{\beta(t)} dw, \quad (6)$$

where $\beta(t) : \mathbb{R} \rightarrow \mathbb{R} > 0$ is the noise schedule of the process, typically taken to be a monotonically increasing linear function of t , and w is the standard Wiener process. The data distribution is defined when $t = 0$, i.e., $x(0) \sim p_{data}$, and a simple, tractable distribution (e.g., Gaussian) is assumed when $t = T$, i.e., $x(T) \sim \mathcal{N}(0, I)$.

With the evolution of the distribution along time t , the neural network accepts both data point x and time t as inputs and produces an estimate of the score function, i.e., $s_\theta(x, t)$. In this case, the denoising score-matching loss function can be written as

$$\text{loss}(\theta) = \mathbb{E}_{t \sim U[0, T]} \mathbb{E}_{p(x(0))} \mathbb{E}_{p(x(t)|x(0))} \left(\lambda(t) \|s_\theta(x(t), t) - \nabla_x \log p(x(t)|x(0))\|_2^2 \right), \quad (7)$$

where $\lambda(t)$ is a positive weighting function and $p(x(t)|x(0))$ can be analytically calculated using (6) to be:

$$p(x(t)|x(0)) \sim \mathcal{N}(\mu(t), \sigma^2(t)I), \quad (8)$$

where

$$\begin{aligned} \mu(t) &= x(0) \exp \left(-\frac{1}{2} \int_0^t \beta(s) ds \right), \\ \sigma^2(t) &= 1 - \exp \left(- \int_0^t \beta(s) ds \right). \end{aligned}$$

The expectations are approximated by averaging over the data points and the integrals can be solved analytically for the specified functions (full implementation details can be found in Table I).

III. APPROXIMATE EM WITH SCORE-BASED PRIOR

The EM algorithm finds a (local) maximum of the posterior distribution (the MAP estimator) by iteratively applying an expectation step (E-step), followed by a maximization step (M-step). The E-step calculates the expectation of the log-likelihood with respect to all admissible configurations of nuisance parameters (translations and rotations in our case). The M-step then finds the parameters maximizing the function found in the E-step, possibly with the addition of a prior term [19].

Direct application of the EM algorithm to the MTD problem is computationally intractable, due to the large number of possible translations that grows quickly with the measurement size. Specifically, under the separation constraint (2), the runtime of a single iteration grows as $\mathcal{O}(N^4)$, where N^2 is the size of the measurement M (1), assuming the measurement is much larger than the target image. Approximate EM that has been introduced in [3] runs in $\mathcal{O}(N^2)$ per iteration, by maximizing an approximation of the log-likelihood function. The approximate EM algorithm has been applied to the two-dimensional MTD problem in [18] and to the cryo-EM problem in [33]. However, in both cases, the prior on the signal has been ignored, which is equivalent to considering an uninformative prior. The main goal of this work is to introduce a data-driven prior into the approximate EM algorithm and to pave the way to learn and implement such prior for cryo-EM datasets.

Approximate EM begins with partitioning the measurement M into $D = (N/L)^2$ non-overlapping patches, each of size $L \times L$ (we assume for convenience that N/L is an integer). The separation condition implies that each patch $M_m, m = 0, \dots, D-1$, can contain either no target image, a full target image, or a part of the target image; overall there are $(2L)^2$ possibilities, excluding rotations. In particular, each patch can be modeled by

$$M_m = CT_{\vec{\ell}_m} ZF_{\phi_m} + \varepsilon_m, \quad \varepsilon_m \sim \mathcal{N}(0, \sigma^2 I), \quad (9)$$

where F is a square image of size $L \times L$, representing a discrete copy of target image f , the operator Z zero-pads L entries to the right and to the bottom of a rotated copy of F , and $T_{\vec{\ell}_m}$ circularly shifts the zero-padded image by $\vec{\ell}_m = (\ell_{mx}, \ell_{my}) \in \mathbb{L} \triangleq \{0, 1, \dots, 2L-1\}^2$ positions, that is,

$$\begin{aligned} (T_{\vec{\ell}_m} ZF_{\phi_m})[i, j] = \\ (ZF_{\phi_m})[(i + \ell_{mx}) \bmod 2L, (j + \ell_{my}) \bmod 2L]. \end{aligned} \quad (10)$$

The operator C then crops the image to size $L \times L$, i.e., $CT_{\vec{\ell}_m} ZF_{\phi_m} = T_{\vec{\ell}_m} ZF_{\phi_m}[0:L-1, 0:L-1]$, and the result is further corrupted by additive white Gaussian noise. Figure 2 provides an illustration of the model. In addition, since the EM algorithm assigns probabilities to rotations (in the E-step), we uniformly discretize the search space of rotations:

$$\phi_m \in \Phi \triangleq \left\{ k \frac{2\pi}{K} \right\}, \quad k = 0, \dots, K-1, \quad (11)$$

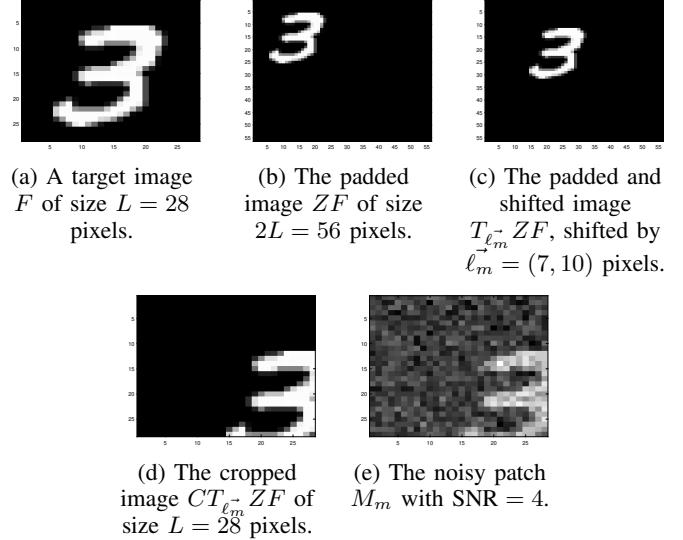


Fig. 2: A step-by-step illustration of the patch generation model described in (9).

where K is a hyper-parameter for the algorithm.

The patches $\{M_m\}_{m=0}^{D-1}$ are our observations, whereas the rotations $\{\vec{\ell}_m\}_{m=0}^{D-1}$ and translations $\{\phi_m\}_{m=0}^{D-1}$ are the unobserved nuisance parameters. Our goal is thus to find the pixel intensity values of F , and we denote by $F^{(\tau)}$ the estimated image at iteration τ of the EM procedure. In the E-step, we calculate the expectation of log-likelihood function of the model, with respect to the conditional distribution of the nuisance parameters given the patches and the current estimate $F^{(\tau)}$:

$$Q(F|F^{(\tau)}) \triangleq \mathbb{E} \left[\log \mathcal{L} | M_0, \dots, M_{D-1}; F^{(\tau)} \right], \quad (12)$$

where \mathcal{L} is the likelihood function, defined as

$$\begin{aligned} \mathcal{L} \triangleq p(M_0, \dots, M_{D-1}, \vec{\ell}_0, \dots, \vec{\ell}_{D-1}, \phi_0, \dots, \phi_{D-1}; F) \\ \approx \prod_{m=0}^{D-1} p(M_m, \vec{\ell}_m, \phi_m; F), \end{aligned} \quad (13)$$

where we neglect statistical dependencies between patches.

We note that Bayes' rule implies

$$p(\vec{\ell}_m, \phi_m | M_m, F^{(\tau)}) = \frac{p(M_m | \vec{\ell}_m, \phi_m, F^{(\tau)}) p(\vec{\ell}_m, \phi_m | F^{(\tau)})}{\sum_{\vec{\ell}' \in \mathbb{L}} \sum_{\phi' \in \Phi} p(M_m | \vec{\ell}', \phi', F^{(\tau)}) p(\vec{\ell}', \phi' | F^{(\tau)})}, \quad (14)$$

where from (9) we can write

$$p(M_m | \vec{\ell}_m, \phi_m, F) \propto \exp \left(-\frac{\|M_m - CT_{\vec{\ell}_m} ZF_{\phi_m}\|_F^2}{2\sigma^2} \right), \quad (15)$$

with the normalization $\sum_{\vec{\ell} \in \mathbb{L}} \sum_{\phi \in \Phi} p(M_m | \vec{\ell}, \phi, F) = 1$. We assume that in each patch $p(\vec{\ell}_m, \phi_m | F) = p(\vec{\ell}_m) p(\phi_m) = \frac{1}{(2L)^2} \cdot \frac{1}{K}$, namely, $\vec{\ell}_m$ and ϕ_m are independent of F and of

Algorithm 1: Approximate EM with score-based prior

Input: measurement M ; noise variance σ^2 ; initial guess $F^{(0)}$; rotation range K (11); max EM iterations \mathcal{T} ; stopping parameter ϵ ; learning rate μ ; regularization weights γ ; prior model $s_\theta(\cdot)$; gradient ascent steps N

Output: an estimate of F

```

1 for  $\tau \leftarrow 0$  to  $\mathcal{T}$  do
2   calculate  $\nabla_F Q(F|F^{(\tau)})$  according to (16)
3   set  $F_0 \leftarrow F^{(\tau)}$ 
4   for  $i \leftarrow 0$  to  $N$  do
5     set
6        $F_{i+1} \leftarrow F_i + \mu (\nabla_F Q(F_i|F^{(\tau)}) + \gamma^{(\tau)} s_\theta(F_i))$ 
7     set  $F^{(\tau+1)} \leftarrow F_N$ 
8     if  $\|F^{(\tau+1)} - F^{(\tau)}\|_F < \epsilon$  then
9       break
9 return  $F^{(\tau)}$ 

```

each other, and are uniformly distributed. Finally, utilizing the approximation in (13), we can write the expected log-likelihood function, up to a constant, as:

$$Q(F|F^{(\tau)}) = \sum_{m=0}^{D-1} \sum_{\vec{\ell} \in \mathbb{L}} \sum_{\phi \in \Phi} p(\vec{\ell}_m, \phi_m | M_m, F^{(\tau)}) \times \log p(M_m | \vec{\ell}_m, \phi_m, F). \quad (16)$$

The M-step updates the current image estimate $F^{(\tau)}$ by maximizing the Q-function (16), with the addition of a regularized prior term:

$$F^{(\tau+1)} = \operatorname{argmax}_F \left(Q(F|F^{(\tau)}) + \gamma^{(\tau)} \log p(F) \right), \quad (17)$$

where $\gamma^{(\tau)}$ is a weighting function, applied to the prior term at iteration τ . In our numerical experiments, the maximization process is being conducted using a fixed number of gradient ascent steps with a fixed learning rate μ and the initial guess of $F^{(\tau)}$. The gradient of the Q-function can be analytically computed, whereas the gradient of the prior term is approximated by our diffusion model $s_\theta(\cdot)$. Thus, a single gradient ascent step can be written as:

$$\begin{aligned} F^{(\tau+1)} &= F^{(\tau)} + \mu \nabla_F \left(Q(F|F^{(\tau)}) + \gamma^{(\tau)} \log p(F) \right) \Big|_{F=F^{(\tau)}} \\ &= F^{(\tau)} + \mu \left(\nabla_F Q(F|F^{(\tau)})|_{F=F^{(\tau)}} + \gamma^{(\tau)} s_\theta(F^{(\tau)}) \right). \end{aligned} \quad (18)$$

The algorithm is summarized in Algorithm 1.

IV. EXPERIMENTAL RESULTS

In this section, we implement our proposed algorithm on two different datasets: (A) Gaussian 5×5 images ($L = 5$) and (B) MNIST 28×28 images ($L = 28$). As our baseline, we consider the plain approximate EM framework (i.e.,

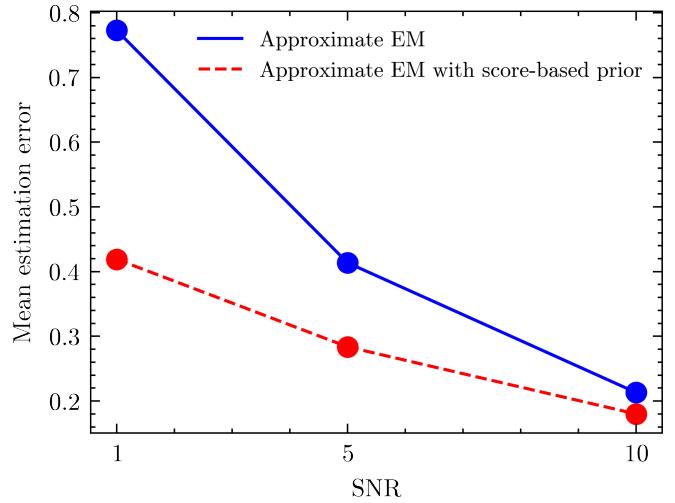


Fig. 3: The mean estimation error of recovering target images F drawn from a Gaussian distribution, as a function of SNR, using approximate EM with and without score-based prior. Evidently, the score-based prior addition is beneficial in all SNRs, especially in low SNR regions.

without diffusion priors) that has been proposed in [18]. We examine both algorithms on a range of SNR values scaled for each dataset; for example, the sparsity of the MNIST images requires lower SNR compared to Gaussian images so that the target images are not easily detected (as in panel (c) of Figure 1). Specifically, for the Gaussian case we test $\text{SNRs} = \{1, 5, 10\}$, and for the MNIST case we test $\text{SNRs} = \{0.1, 0.5, 1\}$. We generate the measurements according to the MTD model (1), with $N = 11L$ and density $\frac{pL^2}{N^2} = 0.1$. The rotations are uniformly sampled from the set $\Phi \triangleq \{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\}$, so that $K = 4$ (11). Finally, for each dataset we examine 10 target images (that are hidden from the prior models during training) and report the mean estimation error, defined by $E = \min_{\phi \in \Phi} \frac{\|F^* - F_\phi\|_2}{\|F^*\|_2}$, where F^* denotes the ground truth image and F_ϕ denotes the estimated image, rotated by ϕ . The error is minimized over the set of rotations because of the intrinsic rotational symmetry of the problem. Additional implementation details are provided in Table I. The code to reproduce the experiments is available at https://github.com/zabatani/2d-mtd_diffusion.

A. Gaussian Images

The score network for this dataset consists of three fully-connected layers with Softplus activation functions. The network has been trained on 10,000 samples drawn from a normal distribution with covariance identity and mean μ , where each entry of μ has been drawn i.i.d. from a uniform distribution over $\{0, 1, 2, 3, 4\}$. The target images were drawn from the same distribution. Training took approximately 10 minutes on an M1 Macbook Air. The results are illustrated in Figure 3. Evidently, our algorithm is superior in all SNRs and exhibits

a significant improvement in recovery accuracy, especially in low SNR levels.

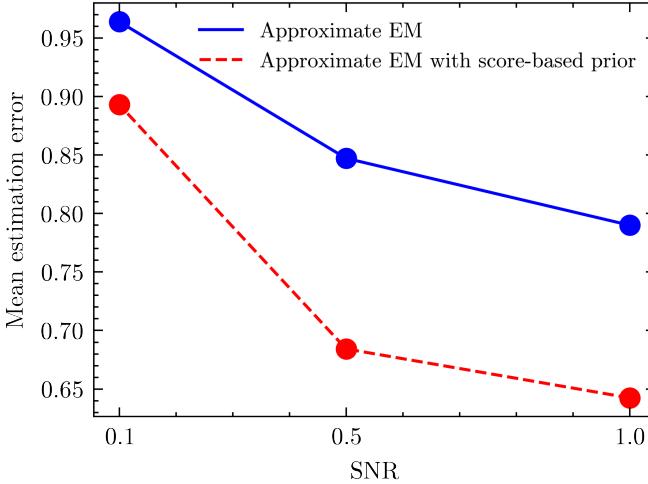


Fig. 4: The mean estimation error of recovering target images F drawn from the MNIST dataset, as a function of SNR, using approximate EM with and without score-based prior. Evidently, the score-based prior addition is noticeably beneficial in all SNRs.

B. MNIST

The score network for this dataset consists of a U-Net architecture [34], with ten convolutional layers and LogSigmoid activation functions. The network has been trained on the entire MNIST dataset (60,000 samples), excluding the 10 target images that have been tested. The training took approximately 10 hours on a T4 GPU machine. In this case, we have noticed that proper weighting of the prior term during gradient ascent plays a key role in our algorithm's success and that a monotonically increasing linear function of τ (the iteration of the EM procedure) works best; see Table I for details. The results are illustrated in Figure 4. Our algorithm is superior in all SNR levels. In addition, we show reconstructions of all digits for a fixed $\text{SNR} = 1$, using both algorithms, in Figure 5. The improvement in recovery quality is visible for all digits.

V. CONCLUSION

In this paper, we proposed to integrate score-based diffusion priors into the approximate EM framework in order to estimate images in the two-dimensional MTD problem. Our method shows a superior image recovery quality compared to [18], which considers an uninformative prior, in various SNRs and on two datasets: a simple Gaussian case, and the more intricate data distribution of MNIST.

This paper is part of an ongoing effort to recover small molecular structures using cryo-EM. The MTD model (1) is a simplified model of the cryo-EM model that includes 3-D rotations (rather than 2-D rotations in (1)) and tomographic projections [10], [11]. The ultimate goal of this research is to

learn a prior of protein structures that were resolved using electron microscopes, and then incorporate it into a computational method for recovering small molecular structures [15], [33]. To learn the priors, we intend to devise 3-D steerable CNN architectures [35] and train them in a denoising score-matching fashion (7). This would require collecting enough data to train such neural networks. To this end, we plan to use databases of resolved protein structures, such as the electron microscopy data bank (EMDB) [36], as well to augment data generated by computational tools, such as AlphaFold [37], [38]. Moreover, we intend to explore additional ways to incorporate the score function into the EM framework that might take advantage of the reverse-time SDE (6) (commonly used for sample generation), for example, by a clever integration between the SDE time evolution and the EM iteration progression. Finally, we intend to integrate our prior into additional algorithms for MTD, such as autocorrelation analysis [1], [2], [15], [39].

ACKNOWLEDGMENT

The research is supported in part by the BSF grant no. 2020159, the NSF-BSF grant no. 2019752, and the ISF grant no. 1924/21.

REFERENCES

- [1] T. Bendory, T.-Y. Lan, N. F. Marshall, I. Rukshin, and A. Singer, "Multi-target detection with rotations," *Inverse Problems & Imaging*, vol. 17, no. 2, 2023.
- [2] N. F. Marshall, T.-Y. Lan, T. Bendory, and A. Singer, "Image recovery from rotational and translational invariants," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 5780–5784.
- [3] T.-Y. Lan, T. Bendory, N. Boumal, and A. Singer, "Multi-target detection with an arbitrary spacing distribution," *IEEE Transactions on Signal Processing*, vol. 68, pp. 1589–1601, 2020.
- [4] S. Kreymer and T. Bendory, "Two-dimensional multi-target detection: An autocorrelation analysis approach," *IEEE Transactions on Signal Processing*, vol. 70, pp. 835–849, 2022.
- [5] Y. LeCun, "The mnist database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998.
- [6] R. Henderson, "The potential and limitations of neutrons, electrons and X-rays for atomic resolution microscopy of unstained biological molecules," *Quarterly Reviews of Biophysics*, vol. 28, no. 2, pp. 171–193, 1995.
- [7] E. Nogales, "The development of cryo-EM into a mainstream structural biology technique," *Nature Methods*, vol. 13, no. 1, pp. 24–27, 2016.
- [8] X.-C. Bai, G. McMullan, and S. H. Scheres, "How cryo-EM is revolutionizing structural biology," *Trends in Biochemical Sciences*, vol. 40, no. 1, pp. 49–57, 2015.
- [9] J. Frank, *Three-dimensional electron microscopy of macromolecular assemblies: visualization of biological molecules in their native state*. Oxford University Press, 2006.
- [10] T. Bendory, A. Bartesaghi, and A. Singer, "Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities," *IEEE Signal Processing Magazine*, vol. 37, no. 2, pp. 58–76, 2020.
- [11] A. Singer and F. J. Sigworth, "Computational methods for single-particle electron cryomicroscopy," *Annual Review of Biomedical Data Science*, vol. 3, pp. 163–190, 2020.
- [12] S. H. Scheres, "RELION: implementation of a Bayesian approach to cryo-EM structure determination," *Journal of Structural Biology*, vol. 180, no. 3, pp. 519–530, 2012.
- [13] A. Punjani, J. L. Rubinstein, D. J. Fleet, and M. A. Brubaker, "cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination," *Nature Methods*, vol. 14, no. 3, pp. 290–296, 2017.
- [14] C. Aguerrebere, M. Delbracio, A. Bartesaghi, and G. Sapiro, "Fundamental limits in multi-image alignment," *IEEE Transactions on Signal Processing*, vol. 64, no. 21, pp. 5707–5722, 2016.

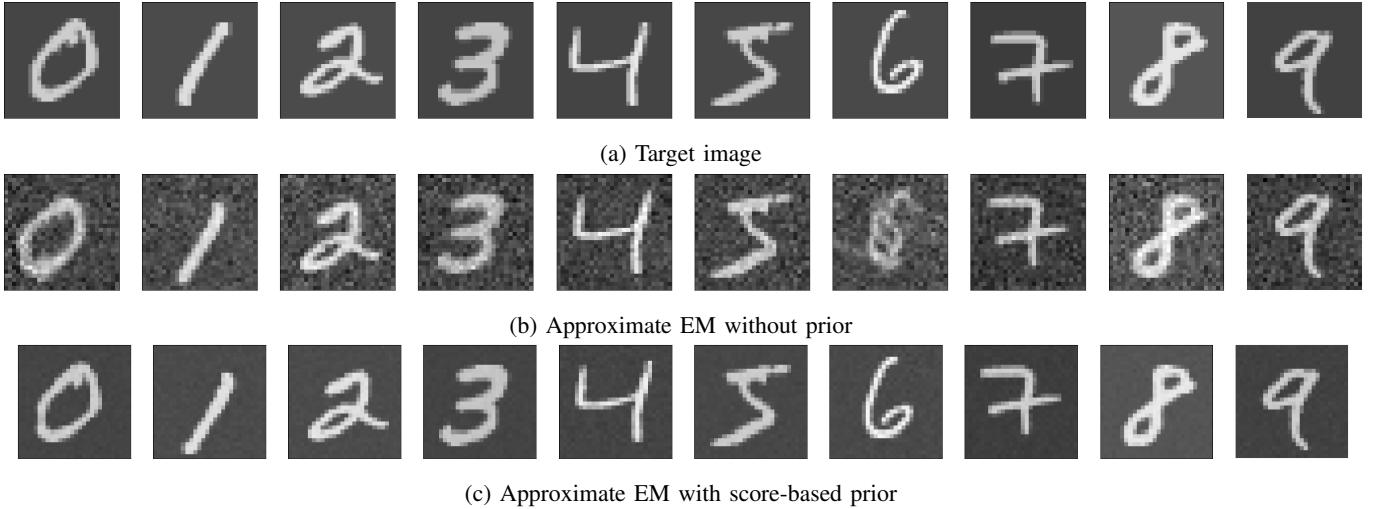


Fig. 5: Recovery of all ten digits from the MNIST dataset of the MTD model with $\text{SNR} = 1$, using the approximate EM algorithm with and without score-based prior. The recovery quality boost achieved by our algorithm (panel (c)) is clearly visible. Besides the background denoising and digit sharpening, it also successfully recovers the target image in cases where the competitor deteriorates (e.g., digit six).

TABLE I: The parameters used for the numerical experiments.

Parameter	Description	Gaussian	MNIST
\mathcal{T}	EM iterations	100	100
N	Gradient ascent steps	1	1
$\gamma^{(\tau)}$	Regularization weighting function (18)	1	τ/\mathcal{T}
μ	Learning rate (18)	$1e-3$	$1e-3$
ϵ	Stopping parameter	$1e-5$	$1e-5$
$F^{(0)}$	Initialization	$U[0, 1]$	Average of ten digits
$\text{loss}(\theta)$	Prior loss function	(5)	(7)
$\beta(t)$	Noise schedule (6)	N/A	$0.1 + (20 - 0.1)t$
$\lambda(t)$	Weighting function (7)	N/A	$1 - e^{\int_0^t \beta(s)ds}$

- [15] T. Bendory, N. Boumal, W. Leeb, E. Levin, and A. Singer, “Toward single particle reconstruction without particle picking: Breaking the detection limit,” *SIAM Journal on Imaging Sciences*, vol. 16, no. 2, pp. 886–910, 2023.
- [16] —, “Multi-target detection with application to cryo-electron microscopy,” *Inverse Problems*, vol. 35, no. 10, p. 104003, 2019.
- [17] E. D’Imprima and W. Kühlbrandt, “Current limitations to high-resolution structure determination by single-particle cryoEM,” *Quarterly Reviews of Biophysics*, vol. 54, 2021.
- [18] S. Kreymer, A. Singer, and T. Bendory, “An approximate expectation–maximization for two-dimensional multi-target detection,” *IEEE Signal Processing Letters*, vol. 29, pp. 1087–1091, 2022.
- [19] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the EM algorithm,” *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, pp. 1–22, 1977.
- [20] D. Kimanius, G. Zickert, T. Nakane, J. Adler, S. Lunz, C.-B. Schönlieb, O. Öktem, and S. H. Scheres, “Exploiting prior knowledge about biological macromolecules in cryo-EM structure determination,” *IUCrJ*, vol. 8, no. 1, pp. 60–75, 2021.
- [21] D. Kimanius, K. Jamali, M. E. Wilkinson, S. Lövestam, V. Velazhahan, T. Nakane, and S. H. Scheres, “Data-driven regularisation lowers the size barrier of cryo-EM structure determination,” *bioRxiv*, pp. 2023–10, 2023.
- [22] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *International conference on machine learning*. PMLR, 2015, pp. 2256–2265.
- [23] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [24] Y. Song and S. Ermon, “Generative modeling by estimating gradients of the data distribution,” *Advances in neural information processing systems*, vol. 32, 2019.
- [25] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” *arXiv preprint arXiv:2011.13456*, 2020.
- [26] P. Dhariwal and A. Nichol, “Diffusion models beat gans on image synthesis,” *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021.
- [27] T. Amit, T. Shaharbany, E. Nachmani, and L. Wolf, “Segdiff: Image segmentation with diffusion probabilistic models,” *arXiv preprint arXiv:2112.00390*, 2021.
- [28] J. Austin, D. D. Johnson, J. Ho, D. Tarlow, and R. Van Den Berg, “Structured denoising diffusion models in discrete state-spaces,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 17981–17993, 2021.
- [29] J. M. L. Alcaraz and N. Strodthoff, “Diffusion-based time series imputation and forecasting with structured state space models,” *arXiv preprint arXiv:2208.09399*, 2022.
- [30] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” *Advances in neural information processing systems*, vol. 27, 2014.
- [31] Q. Liu, J. Lee, and M. Jordan, “A kernelized Stein discrepancy for goodness-of-fit tests,” in *International conference on machine learning*. PMLR, 2016, pp. 276–284.
- [32] A. Hyvärinen and P. Dayan, “Estimation of non-normalized statistical models by score matching,” *Journal of Machine Learning Research*, vol. 6, no. 4, 2005.

- [33] S. Kreymer, A. Singer, and T. Bendory, “A stochastic approximate expectation-maximization for structure determination directly from cryo-EM micrographs,” *arXiv preprint arXiv:2303.02157*, 2023.
- [34] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*. Springer, 2015, pp. 234–241.
- [35] M. Weiler, M. Geiger, M. Welling, W. Boomsma, and T. S. Cohen, “3d steerable cnns: Learning rotationally equivariant features in volumetric data,” *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [36] C. L. Lawson, A. Patwardhan, M. L. Baker, C. Hryc, E. S. Garcia, B. P. Hudson, I. Lagerstedt, S. J. Lutcke, G. Pintilie, R. Sala *et al.*, “EMDataBank unified data resource for 3DEM,” *Nucleic acids research*, vol. 44, no. D1, pp. D396–D403, 2016.
- [37] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko *et al.*, “Highly accurate protein structure prediction with AlphaFold,” *Nature*, vol. 596, no. 7873, pp. 583–589, 2021.
- [38] T. C. Terwilliger, D. Liebschner, T. I. Croll, C. J. Williams, A. J. McCoy, B. K. Poon, P. V. Afonine, R. D. Oeffner, J. S. Richardson, R. J. Read *et al.*, “AlphaFold predictions are valuable hypotheses and accelerate but do not replace experimental structure determination,” *Nature Methods*, pp. 1–7, 2023.
- [39] Z. Kam, “The reconstruction of structure from electron micrographs of randomly oriented particles,” *Journal of Theoretical Biology*, vol. 82, no. 1, pp. 15–39, 1980.