

# A Molecular Prior Distribution for Bayesian Inference Based on Wilson Statistics

Marc Aurèle Gilles<sup>\*†</sup>, Amit Singer<sup>‡</sup>

May 3, 2022

## Abstract

**Background and Objective:** Wilson statistics describe well the power spectrum of proteins at high frequencies. Therefore, it has found several applications in structural biology, e.g., it is the basis for sharpening steps used in cryogenic electron microscopy (cryo-EM). A recent paper gave the first rigorous proof of Wilson statistics based on a formalism of Wilson’s original argument. This new analysis also leads to statistical estimates of the scattering potential of proteins that reveal a correlation between neighboring Fourier coefficients. Here we exploit these estimates to craft a novel prior that can be used for Bayesian inference of molecular structures.

**Methods:** We describe the properties of the prior and the computation of its hyperparameters. We then evaluate the prior on two synthetic linear inverse problems, and compare against a popular prior in cryo-EM reconstruction at a range of SNRs.

**Results:** We show that the new prior effectively suppresses noise and fills-in low SNR regions in the spectral domain. Furthermore, it improves the resolution of estimates on the problems considered for a wide range of SNR and produces Fourier Shell Correlation curves that are insensitive to masking effects.

**Conclusions:** We analyze the assumptions in the model, discuss relations to other regularization strategies, and postulate on potential implications for structure determination in cryo-EM.

## 1 Introduction

In his seminal 1942 paper, Arthur Wilson showed that the power spectrum of molecules is approximately flat at high frequencies [23]. The core assumption Wilson made is that atoms act as if they were randomly positioned at high frequency. Wilson used that randomness assumption and the observation that atom positions map to waves in the spectral domain to argue that the cross-terms of the waves sum up incoherently, implying a flat power spectrum. Wilson statistics finds several applications in structural biology. For example, in the field of cryogenic electron microscopy (cryo-EM), it is the basis of B-factor sharpening [15], a post-processing step aimed at increasing the contrast of experimentally obtained reconstructions.

A recent paper [18] provided a rigorous derivation of Wilson statistics based on a formalism of Wilson’s argument. Using that formalism, the author derived other forms of statistics, particularly a mean and covariance estimate for the scattering potential of molecules. In the present paper, we develop a prior based on these estimates and study their use in Bayesian inference of protein structure in cryo-EM. After we state background on the use of maximum a posteriori estimation and priors in cryo-EM in section 2.1, we describe the new prior in section 2.2. Then, in section 3, we evaluate the prior on synthetic linear inverse problems and discuss its properties and further applications in section 4.

The code to reproduce numerical experiments appearing in this manuscript is publicly available at: [https://github.com/ma-gilles/wilson\\_prior](https://github.com/ma-gilles/wilson_prior).

## 2 Background

### 2.1 Priors in cryo-EM

In standard single particle cryo-EM structure determination, the 3D scattering potential map of a molecule  $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}$  is determined from particle projection images  $y_i : \mathbb{R}^2 \rightarrow \mathbb{R}$ . The critical computational step for reconstruction is the iterative

<sup>\*</sup>A.S. and M.A.G are supported in part by AFOSR FA9550-20-1-0266, the Simons Foundation Math+X Investigator Award, NSF BIGDATA Award IIS-1837992, NSF DMS-2009753, and NIH/NIGMS 1R01GM136780-01

<sup>†</sup>Program in Applied and Computational Mathematics, Princeton University, Fine Hall, Washington Road, Princeton, NJ 08544-1000, mg6942@princeton.edu

<sup>‡</sup>Department of Mathematics and PACM, Princeton University, Fine Hall, Washington Road, Princeton, NJ 08544-1000

refinement, typically formulated as maximum a posteriori estimation. In that step, the reconstructed molecule is the one that maximizes the posterior probability:

$$\phi_{\text{MAP}} := \arg \max_{\phi} p(\phi|y) \stackrel{\text{Bayes' law}}{=} \arg \max_{\phi} p(y|\phi)p(\phi) , \quad (1)$$

where  $p(y|\phi)$  is the likelihood function (the conditional probability of the images given the molecule), and  $p(\phi)$  is the prior distribution over molecules. The prior encodes our belief about the distribution of molecules before any observation and often imposes particular properties on the inferred scattering potential (e.g., smoothness). Priors are a form of regularization: they add information to solve an ill-posed problem and avoid overfitting.

In many inverse problems, including cryo-EM reconstructions, the quality of the inferred estimate depends heavily on the choice of prior. Thus, the cryo-EM community has given much attention to crafting appropriate priors. One of the earliest software packages to use MAP estimation for 3D reconstruction in cryo-EM was RELION [16]. The initial version of RELION used a Gaussian prior distribution where each frequency is independent with zero mean and variance equal to the spherically-averaged power spectrum of the molecule at that frequency. We call this prior the diagonal prior, denoted as:

$$p_{\text{diagonal}}(\hat{\phi}) := \mathcal{N}(0, D_{\text{PS}(\phi)}) , \quad (2)$$

where  $\hat{\phi}$  denotes the Fourier transform of  $\phi$ ,  $\mathcal{N}(\mu, \Sigma)$  denotes a normal distribution with mean  $\mu$  and covariance  $\Sigma$ ,  $D_v$  denotes a diagonal matrix with diagonal  $v$ , and  $\text{PS}(\phi)(\xi) := \int_{S^2} |\hat{\phi}(\|\xi\|\omega)|^2 d\omega$  denotes the spherically-averaged power spectrum of  $\phi$ . The diagonal prior has several attractive properties: it enforces smoothness on the reconstructed molecule, and it is computationally cheap thanks to the independence assumption of Fourier components.

The initial version of cryoSPARC [10], another popular software for cryo-EM reconstruction, used a spatially-independent exponential (SIE) prior:

$$p_{\text{SIE}}(\phi) := \text{Exp}(\phi) \quad (3)$$

where  $\text{Exp}(\mu)$  denotes the exponential distribution with mean  $\mu$ . The advantages of this prior are that it imposes positivity on the scattering potential, and it is computationally convenient thanks to the independence of voxel values. We note that despite the statistical interpretation of priors in eq. (1), both the diagonal and SIE priors are chosen out of mathematical and computational convenience rather than a belief about the distribution of molecules. This may bias the reconstruction process [3]. In newer versions of RELION and cryoSPARC, both software packages use a modified version of the diagonal prior as defined above. In that prior, each Fourier frequency is modeled as independent Gaussian with variance set proportionally to the SNR as estimated by Fourier Shell Correlation (FSC) between half maps [17, 11].

In recent years, *implicit* regularization schemes have become popular in the computational imaging community, where there is often no explicit prior function  $p(\phi)$ . Instead, an operator mimics the regularizer’s action in an iterative algorithm (e.g., the operator may act as a gradient descent step [14, 13] or a proximal operator [20, 5]). One example in cryo-EM, implemented in RELION, uses a denoising neural network to regularize expectation-maximization iterations [8]. Another notable example, implemented in the non-uniform refinement option of cryoSPARC, regularizes the iteration using a smoothing kernel whose parameters are set adaptively by cross-validation [12]. Implicit regularization schemes are more general than regularization using an explicit prior<sup>1</sup> and can leverage powerful machine learning techniques that sometimes yield impressive results, e.g. [8, 12] both report that the prior nearly halves the attained resolution on some test cases compared to the traditional prior. On the other hand, implicit regularization schemes lose most theoretical guarantees and statistical interpretation granted by Bayesian inference, sometimes leading to overfitting. Overfitting is particularly problematic in cryo-EM, where ground truth about the 3D structure is often unavailable, making diagnosing overfitting difficult. In the rest of this paper, we focus on building an explicit prior, similar to the ones in eqs. (2) and (3) but derived from Wilson statistics.

## 2.2 The Wilson prior

The formalism used to derive Wilson statistics in [18] is the random “bag of atoms” model. In that model, a molecule consists of a collection of  $N_{\text{atoms}}$  atoms whose positions  $X_1, X_2, \dots, X_{N_{\text{atoms}}}$  are independently and identically distributed (i.i.d.) with probability density function  $g : \mathbb{R}^3 \rightarrow \mathbb{R}^+$  that models the molecule’s shape. The scattering potential of a molecule  $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}$  is modeled as

$$\phi(x) = \sum_{i=1}^{N_{\text{atoms}}} f(x - X_i) \quad (4)$$

where  $f$  is the scattering potential of a single atom. For simplicity, here we assume that the atoms are identical and revisit to this assumption in section 4. If  $f$ ,  $g$ , and  $N_{\text{atoms}}$  are fixed, eq. (4) implicitly defines a probability distribution over molecules that can, in principle, be used as a prior for Bayesian inference. Unfortunately, there is no easy formula for the probability density function of that distribution, so it is impractical to do so. Instead, we use the first two moments of the bag of atoms

<sup>1</sup>E.g., the impossibility result in [13] describes the conditions under which an implicit denoising regularizer cannot be written as an explicit regularizer.

model to craft a Gaussian prior that resembles the original distribution but is easier to compute. We express this prior in the Fourier domain, but an equivalent formulation in the spatial domain is described in section 6.1. The first two moments of the distribution of the Fourier transform of  $\phi$  are derived in [18]:

$$\hat{\mu}_{\text{Wilson}}(\xi) := \mathbb{E}[\hat{\phi}(\xi)] = N_{\text{atoms}} \hat{f}(\xi) \hat{g}(\xi) , \quad (5)$$

$$\hat{\Sigma}_{\text{Wilson}}(\xi_1, \xi_2) := \mathbb{E}[(\hat{\phi}(\xi_1) - \hat{\mu}_{\text{Wilson}}(\xi_1))(\hat{\phi}(\xi_2) - \hat{\mu}_{\text{Wilson}}(\xi_2))] = N_{\text{atoms}} \hat{f}(\xi_1) \overline{\hat{f}(\xi_2)} (\hat{g}(\xi_1 - \xi_2) - \hat{g}(\xi_1) \overline{\hat{g}(\xi_2)}) . \quad (6)$$

where  $\hat{f}$ ,  $\hat{g}$  denote the Fourier transforms of  $f$  and  $g$  respectively. We thus define the Wilson prior as:

$$p_{\text{Wilson}}(\hat{\phi}) = N(\hat{\mu}_{\text{Wilson}}, \hat{\Sigma}_{\text{Wilson}}) . \quad (7)$$

Due to the non-Gaussianity of the bag of atoms model, the Wilson prior and the bag of atoms model are not the same. In particular, negative values in the spatial domain are possible in the Wilson prior but not in the bag of atoms model (assuming  $f(x) \geq 0$ ). Nevertheless, the Wilson prior is the maximum entropy distribution, which matches the first and second-order statistics derived from the bag of atoms model and allows us to derive a fast algorithm. Figure 1 illustrates the differences between the bag of atoms model and three priors: Wilson, diagonal, and SIE, as defined in eqs. (2), (3) and (7).

The diagonal and SIE priors represent two extremes: in the diagonal prior, frequencies in the spectral domain are independent, and in the SIE prior, voxels in the spatial domain are independent. The Wilson prior is a middle ground: the decay of the Fourier-transformed shape function  $\hat{g}(\xi)$  dictates the correlation between two frequencies, and the decay of the function  $f(x)$  dictates the correlation between two positions. This latter fact is most evident from the Wilson statistics expressed in the spatial domain (see section 6.1 in the appendix). In fig. 1, observe that projection images generated by the diagonal prior have a similar texture as the ones generated by the Wilson prior, but they do not capture the shape of the molecule. In section 6, we explain this phenomenon by showing that the diagonal and Wilson priors are equal in the limit on a uniform shape function on the entire domain. As a result, we interpret the Wilson prior as a generalization of the diagonal prior beyond the case of a uniform shape function on the entire domain.

If no information about the molecule’s shape is available, an uninformative prior reflecting equal probability of the position of atoms at any point in the domain is appropriate. Thus, the diagonal prior is the logical choice. However, during the iterative refinement of 3D reconstruction, we can often observe the low frequency of the shape function at early iterations. Remarkably, these *low frequencies* capture most of the *correlation of high frequencies*. This is a consequence of a result in [18]: under mild assumptions on the shape function, its Fourier transform decays quadratically with frequency  $|\hat{g}(\xi)| = \mathcal{O}(\xi^{-2})$ . At high frequency, the decay implies that  $|\hat{g}(\xi_1 - \xi_2)| \gg |\hat{g}(\xi_1) \overline{\hat{g}(\xi_2)}|$  therefore the covariance of the Wilson statistics is:

$$\hat{\Sigma}_{\text{Wilson}}(\xi_1, \xi_2) = N_{\text{atoms}} \hat{f}(\xi_1) \overline{\hat{f}(\xi_2)} (\hat{g}(\xi_1 - \xi_2) - \hat{g}(\xi_1) \overline{\hat{g}(\xi_2)}) \approx N_{\text{atoms}} \hat{f}(\xi_1) \overline{\hat{f}(\xi_2)} \hat{g}(\xi_1 - \xi_2) .$$

Since  $\hat{f}(\xi)$  decays slowly in the spectral domain, the quadratic decay of the shape function  $\hat{g}(\xi)$  dictates the correlation of neighboring frequencies. Therefore, Fourier coefficients at high frequencies are correlated to their neighboring Fourier coefficients, but that correlation decays quadratically with the neighborhood size. It follows that capturing  $\hat{g}(\xi)$  for small  $\xi$  (that is, the low frequencies of the shape function) is sufficient to capture most of the correlation between high frequencies. It is this correlation we seek to exploit in developing the Wilson prior.

## 3 Application to inverse problems

### 3.1 Problem statement

We evaluate the Wilson prior on linear inverse problems; that is, by inferring the 3D scattering potential of the molecule  $\phi$  from measurements  $y$ , generated by:

$$y = A\phi + \epsilon , \quad (8)$$

where  $A$  is the measurement matrix, and  $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$  is additive Gaussian white noise. In particular, we study the denoising problem where  $A$  is the identity and the deconvolution problem where  $A$  is diagonal in the Fourier domain. Both problems are significantly simpler than 3D reconstruction of cryo-EM structures, but the expectation step of expectation-maximization takes the form of eq. (8); where the diagonal matrix  $A$  also accounts for marginalizing over the latent variables.

We estimate  $\phi$  from  $y$  with MAP estimation as in eq. (1). When the noise and the prior are Gaussian distributions (as in the case for Wilson and diagonal priors), the optimization problem is rewritten as a least-squares problem:

$$\phi^{MAP} = \arg \min_{\phi} \left\{ \frac{1}{2\sigma^2} \|A\phi - y\|_2^2 + \frac{1}{2} \|\phi - \mu\|_{\Sigma^{-1}}^2 \right\} , \quad (9)$$

where  $\mu, \Sigma$  are the prior mean and prior covariance. The solution of eq. (9) is the result of the Wiener filter:

$$\phi^{MAP} = (I - HA)\mu + Hy, \quad H = \Sigma A^* (A \Sigma A^* + \sigma^2 I)^{-1} . \quad (10)$$




Figure 1: Illustration of different prior distributions and comparison with a reference protein. (a) Projection image of the SARS-CoV-2 spike glycoprotein (PDB 6VXX), used for reference. (b) Zoomed-in view of the spike protein at the same scale as images sampled from the priors; meant to illustrate the texture of real proteins at this scale. (c) Random samples from four distributions: the bag of atoms model, the Wilson prior, the diagonal prior, and the spatially-independent exponential prior. (1): the bag of atoms model with  $N_{\text{atoms}} = 8871$ , the shape function  $g$  is a uniform distribution over the ball of radius  $40 \text{ \AA}$  and the atom scattering function  $f$  is the carbon atom shape described in [9]. (2): the Wilson prior with the same  $N_{\text{atoms}}$ ,  $g$  and  $f$ . (3): the diagonal prior where the variance of the frequencies is set to the expected power spectrum of the bag of atoms model. (4): the SIE prior with mean equal to the expected value under the bag of atoms model. The Wilson prior’s distribution is close to the bag of atoms model but contains a small number of negative values. The diagonal prior has a similar texture but does not capture the support or voxel distribution. The SIE prior does not capture the texture of projection images due to the independence assumption of voxels in real space. Mismatches between the prior and the actual distribution of molecules can bias the reconstruction by Bayesian inference. We emphasize that both the diagonal and SIE prior are chosen for computational convenience, whereas the Wilson prior was derived from a molecular model.

Plugging in different prior means and covariances result in different MAP estimates, which we refer to as MAP-diagonal and MAP-Wilson. Both the diagonal and the Wilson priors have hyperparameters that depend on the true molecule: the spherically-averaged power spectrum in the case of the diagonal prior, the shape function  $g$ , atomic scattering function  $f$ , and the number of atoms in the case of Wilson prior. In cryo-EM applications, the spherically-averaged power spectrum can often be well approximated from the data before the reconstruction is performed [19]. However, the shape of the molecule is not typically known a priori. We address this in the following section.

### 3.2 Estimating parameters of the Wilson prior

The Wilson prior depends on three quantities: the number of atoms  $N_{\text{atoms}}$ , the atom scattering function  $f(x)$  and the shape function  $g(x)$ . To fix  $N_{\text{atoms}}$  and  $f(x)$ , we assume that the atomic composition of the target molecule is known. We set  $N_{\text{atoms}}$  to the ground truth number of non-hydrogen atoms and set  $f$  equal to the weighted average of the scattering functions of atoms in the molecule:

$$f(x) = \sum_a p_a f_a(x)$$

where  $a$  indexes the atom type,  $p_a$  denotes the proportion of atom  $a$  in the molecule, and  $f_a$  the scattering potential of atom  $a$ , each modeled by a sum of 5 Gaussians as in [9].

To approximate the shape function  $g$ , we assume that we have a (possibly low resolution) estimate of the molecule  $\phi_{\text{guess}}$ . We get a formula for  $g$  by approximating the mean of the Wilson statistics as  $\phi_{\text{guess}}$ :

$$\hat{\phi}_{\text{guess}}(\xi) \approx \hat{\mu}_{\text{Wilson}}(\xi) = N_{\text{atoms}} \hat{f}(\xi) \hat{g}(\xi)$$

Solving for  $g(x)$ , we get:  $g(x) = \mathcal{F}^{-1} \left\{ \frac{\hat{\phi}_{\text{guess}}(\xi)}{N_{\text{atoms}} \hat{f}(\xi)} \right\}$ , i.e.,  $g(x)$  is proportional to our guess image deconvolved by the atom scattering function. If  $\phi_{\text{guess}}$  were precisely the true molecule scattering function and all atom scatterings were equal to  $f(x)$ , then the estimated  $g(x)$  would be a sum of delta function at the true atomic positions. To account for noise and modeling mismatch, we propose to relax this approximation by convolving this quantity with a Gaussian kernel  $G_\nu(x) = \exp(-\|x\|^2/(2\nu^2)) / (2\pi\nu^2)^{3/2}$ . Finally, since  $g(x)$  is a probability distribution of the position of atoms, it must lie in the probability simplex  $\Delta = \{g | \int g(x) dx = 1, g(x) \geq 0\}$ . Numerically, this condition ensures that the covariance matrix in eq. (6) is positive semi-definite. To enforce this constraint, we orthogonally project our guess onto  $\Delta$  (denoted  $\Pi_\Delta$ ). In summary, we estimate the shape function  $g$  from a guess of the molecule  $\phi_{\text{guess}}$  by:

$$g_{\phi_{\text{guess}}, \nu}(x) = \Pi_\Delta \left( G_\nu(x) * \mathcal{F}^{-1} \left\{ \frac{\hat{\phi}_{\text{guess}}(\xi)}{N_{\text{atoms}} \hat{f}(\xi)} \right\} \right). \quad (11)$$

In all numerical experiments, we take  $\phi_{\text{guess}}$  to be the MAP-diagonal estimate and set the kernel width  $\nu = 1\text{\AA}$ . We discuss alternatives to estimate  $g(x)$  in section 4.

We mention that convolving an empirical distribution by a kernel is known as kernel density estimation; a non-parametric method typically used to estimate a random variable’s probability density function based on a finite data sample (e.g., [6]). Here, we use it in a different framework since we never directly observe samples of atomic locations. The effect of kernel density estimation is to smooth the distribution, which is interpreted as regularizing the estimated distribution. This strategy also recalls the regularization scheme of [12], where a non-stationary Butterworth kernel is used to regularize in cryo-EM reconstruction. A significant difference is that we use the kernel to set a hyperparameter of the prior, whereas [12] uses the kernel directly to regularize the molecule.

### 3.3 Evaluation of the Wilson prior on synthetic data

We construct a ground truth scattering potential  $\phi$  from the reported atomic coordinates of the spike protein (PDB 6VXX) by evaluating the following sum on a regular 3D frequency grid using the NUFFT [2, 1]:

$$\hat{\phi}(\xi) = \sum_a \hat{f}_a(\xi) \sum_{i_a} \exp(-2\pi i \langle \xi, x_{i_a} \rangle) \quad (12)$$

where  $a$  is the atom type, and  $\hat{f}_a(\xi)$  is the atomic scattering model reported in [9]. The final ground truth is obtained by inverse discrete Fourier transform and masking using the ground truth mask reported below, which suppresses oscillations caused by truncation of the spectrum.

We report the Fourier Shell Correlation (FSC) between the ground truth and reconstructed molecules, both with and without masks. Masked FSC is the standard method to estimate resolution in cryo-EM, but using an overly tight mask will inflate the FSC, resulting in an over-optimistic resolution estimate. On the other hand, using no masks often underestimates the map’s resolution, as the noise in the background dominates the FSC at higher frequencies. We use the resolution criterion “masked FSC equal to 0.5” since we compare our reconstruction to the ground truth [15]. We use a ground truth

mask generated by the software package EMDA [22] by placing a sphere of radius  $3\text{\AA}$  at each atomic location, followed by dilation and convolution with a Gaussian.

In fig. 2, we compare the result of the denoising Wiener filter using the diagonal and Wilson prior at different SNRs. We define  $\text{SNR} = \frac{\|\phi\|_2^2}{\sigma^2 N}$  where  $N = 301^3$  is the number of voxels. Following the implementation in [16], we scale the covariance matrix of the diagonal prior (but not the Wilson prior) by a constant  $T = 4$ . This scaling factor improves the reconstruction of the MAP-diagonal estimate, as was also empirically observed in [16]. The first noteworthy result is that the FSC of the MAP-Wilson estimate is insensitive to the mask compared to the MAP-diagonal estimate. In section 6.1 in the appendix, we show that if we take the atomic scattering function to be a delta function (which is approximately true for low-resolution problems), the support of the shape function contains the support of the MAP-Wilson estimate. Hence, estimating  $g$  is analogous to calculating a mask, and the MAP-Wilson estimate has a masking effect.

Within the mask, the MAP-Wilson's resolution is significantly higher than the MAP-diagonal:  $2.6\text{\AA}$  vs  $3.1\text{\AA}$  at  $\text{SNR} = 10^{-2}$ , and  $8.1\text{\AA}$  vs  $10.1\text{\AA}$  at  $\text{SNR} = 10^{-3}$ . We interpret that the diagonal prior over-smooths the protein, whereas the Wilson prior exploits the correlation of neighboring Fourier components to suppress more noise while retaining higher frequency




Figure 2: Denoising of the spike protein contaminated with Gaussian white noise at different SNRs. Left: projected images of the noisy observations. Middle: isosurface visualization of the ground truth, the MAP-diagonal estimate, and the MAP-Wilson estimate. Isosurface levels were qualitatively set to reflect the structure best. Right: FSC plots.

In fig. 3, we consider the problem of deconvolution with the contrast transfer function (CTF):

$$\text{CTF}(\xi) = \sin(-\pi\lambda d\|\xi\|^2 + C_s\lambda^3\|\xi\|^4\pi/2 - \alpha) \exp(-B\|\xi\|^2/4) ,$$

with  $\lambda = 2.51$  nm,  $d = 1.5$  nm,  $C_s = 2.0$  nm,  $\alpha = 0.1$ , and  $B = 80 \text{ \AA}^2$ . Note that the zero crossings of the CTF annihilate some of the measured Fourier coefficients; thus, the SNR is small around these frequencies. This produces the visible oscillations in the unmasked FSC of the MAP-diagonal estimate. In contrast, the oscillations are not present in the MAP-Wilson case. Therefore, we interpret that the Wilson prior uses the correlation between neighboring frequencies to fill zero-crossing gaps. The ability to fill low-SNR gaps may be helpful in cryo-EM applications as SNR is non-uniform in Fourier space due to the CTF and the non-uniform distribution of molecule rotations.




Figure 3: Deconvolution of the spike protein contaminated with Gaussian white noise at different SNRs. Left: projected images of the noisy observations. Middle: isosurface visualization of the ground truth, MAP-diagonal estimate, and MAP-Wilson estimate. Right: FSC plots.

### 3.4 Computational matters

The Wiener filter in eq. (10) involves computing the result of

$$Hv = \Sigma A^* (A \Sigma A^* + \sigma^2 I)^{-1} v . \quad (13)$$

$\Sigma_{\text{diagonal}}$  and  $A$  are diagonal matrices in the Fourier basis, therefore the MAP-diagonal estimate can be computed by elementwise multiplication in  $\mathcal{O}(N)$  operations, where  $N$  is the number of voxels. However,  $\Sigma_{\text{Wilson}}$  is dense, so the same strategy does not apply. Instead, we solve the linear system in eq. (13) with the conjugate-gradient algorithm (CG) [7].

The computational bottleneck of CG is the matrix-vector products with the matrix  $A\Sigma A^* + \sigma^2 I$ . Thanks to the special structure of  $\Sigma_{\text{Wilson}}$ <sup>2</sup> these matrix-vector products can be computed in  $\mathcal{O}(N \log(N))$  using the FFT. The computational cost of solving eq. (13) approximately is thus  $\mathcal{O}(\kappa N \log(N))$  where  $\kappa$  is the number of CG iterations. In experiments, we terminate CG when  $\kappa = 100$  or we achieve a tolerance of  $10^{-8}$ . Finally, computation of the shape function as described in eq. (11) costs  $\mathcal{O}(N \log(N))$  operations with the FFT and fast projection algorithms for the probability simplex [21].

We note that although the diagonal prior is faster than the Wilson prior<sup>3</sup>, solving eq. (13) for either prior is a negligible portion of the computation performed in 3D refinement. Indeed, the dominant cost of the maximization step of a typical expectation-maximization iteration used in 3D refinement is marginalizing over latent variables and summing over projections images to form the matrix and the right-hand side, which dwarfs the cost of the CG iteration. The other additional computation is updating the prior mean and covariance (which one could consider part of the E step in an EM algorithm). Since we do not explicitly form the matrix  $\Sigma_{\text{Wilson}}$ , this only requires updating the shape function  $g$ . In this paper, we computed  $g$  as in eq. (11) at negligible computational expense, but more involved strategies such as the ones discussed in section 4 could incur a more substantial cost. We leave this question to future work.

## 4 Discussion

We showed a prior based on Wilson statistics outperforms a diagonal prior at solving two linear inverse problems within the MAP estimation framework. We argue that the Wilson prior encodes more information about molecules, namely, the correlation between neighboring Fourier coefficients. This correlation is used to suppress noise and fill gaps in low SNR regions of the spectral domain.

In contrast with implicit regularization, whose assumptions and convergence properties are opaque, we derived the Wilson prior from first principles. It fits squarely within the MAP framework, and we can guarantee its convergence and analyze violated assumptions. The Wilson prior is the best Gaussian approximation of the bag of atoms model, which makes the following assumptions:

- **Assumption 1** *Atoms all have the same scattering function.*
- **Assumption 2** *Atoms positions are independently identically distributed.*
- **Assumption 3** *Atoms positions are distributed with probability density function  $g(x)$ .*

We made assumption 1 for convenience here, but it is not necessary, and one could extend this framework to multiple atom types. However, numerical experiments indicate that assumption 1 is unimportant as there is a minimal loss of accuracy in denoising real molecules compared to fake molecules modified to have a single atom type.

Assumption 2 is the core assumption made in Wilson statistics, but it is violated in practice. For example, the position of subsequent atoms is fixed in an alpha-helix. Avoiding an independence assumption is difficult when crafting a prior, but considering atom interactions might bring further improvement if feasible.

Assumption 3 can be problematic as a poor choice of  $g(x)$  will lead to bias. For example, we showed that  $g(x)$  acts as a mask in the denoising case, so carefully choosing its support is crucial. We have advocated estimating  $g(x)$  from an estimate of the molecule, but this could lead to overfitting, for the same reason that crafting a mask from an estimate may lead to overfitting [19]. On the other hand, we argued in section 2.2 that only the lowest frequencies of  $g(x)$  need to be accurately computed for most of the prior correlation to be correctly estimated. We also showed in synthetic experiments that a simple choice of  $g(x)$  leads to significant resolution improvement with no visible overfitting. It is unclear whether this result will hold when the prior is used in an expectation-maximization algorithm. In that case, a more robust framework may be to estimate  $g$  as part of the maximum likelihood estimation. The maximum likelihood estimate of  $g$  and  $\phi$  given measurements  $Y$  is obtained by the following constrained minimization problem:

$$\phi_{\text{MLE}}, g_{\text{MLE}} = \arg \min_{g \in \Delta, \phi} \left\{ \frac{1}{2\sigma^2} \|Y - A\phi\|_2^2 + \frac{1}{2} \|N_{\text{atoms}} \hat{f} \hat{g} - \phi\|_{\Sigma_g^{-1}}^2 + \log \det(\Sigma_g) \right\}, \quad (14)$$

$$\text{where } \Sigma_g(\xi_1, \xi_2) = N_{\text{atoms}} \hat{f}(\xi_1) \left( g(\xi_1 - \xi_2) - \hat{g}(\xi_1) \overline{\hat{g}(\xi_2)} \right) \overline{\hat{f}(\xi_2)}. \quad (15)$$

Minimizing this quantity with an alternating scheme (e.g., ADMM [4]) would produce iterations similar to ones outlined above: first minimize over  $\phi$ ; which is equivalent to a Wiener filter as in eq. (10), and then minimize over  $g$ ; corresponding to a non-linear version of the projection presented in section 3.2. Our results suggest that replacing mask estimation (a routine part of 3D reconstruction) with the estimation of the shape function would be beneficial for several reasons:

<sup>2</sup> $\Sigma$  is a diagonally scaled convolution operator plus a rank one matrix. Diagonal scaling costs  $\mathcal{O}(N)$  operations, convolution  $\mathcal{O}(N \log(N))$  and the matrix vector product with a rank-1 matrix is  $\mathcal{O}(N)$ , thus the total cost is  $\mathcal{O}(N \log(N))$ .

<sup>3</sup>The convergence of CG, and therefore the runtime of the Wilson prior, varies with the condition number of the matrix  $(A\Sigma A^* + \sigma^2 I)$ , which depends on the CTF and the noise level. As an example, in our implementation, the denoising Wiener filter takes 26 seconds using the diagonal prior and 6 minutes using the Wilson prior for SNR= $10^{-2}$  and a grid of size  $301^3$ , on a Macbook Pro (16 GB RAM and 2 GHz Quad-Core Intel Core i5).

1. It improves the resolution of the estimates.
2. The corresponding FSC curves are insensitive to ground truth masks.
3. One can estimate the shape function within the Bayesian framework, which is typically not the case for a mask.

To be integrated within a pipeline for 3D refinement, the prior requires two generalizations: a variable noise model and an adaptive variance estimation. The former is a straightforward extension of the framework presented, e.g., the noise model used in RELION can be implemented in eq. (10) by replacing the matrix  $\sigma^2 I$  by a different diagonal matrix. The latter is a crucial property of the default prior used in RELION and cryoSPARC: the diagonal of the prior covariance matrix is updated at each iteration with the SNR as estimated by the FSC between half maps. This strategy is a type of implicit frequency marching [3]: the regularization decreases for the high frequencies at each iteration, gradually increasing the resolution of the estimate. The same behavior does not exist for the prior we presented here since the variance of  $\Sigma_{\text{Wilson}}$  at high frequencies is approximately  $N_{\text{atoms}}|\hat{f}(\xi)|^2$ , which is independent of the current estimate. One way to achieve the same result would be to scale the covariance matrix with the factor computed by FSC. An alternative would be to treat the refinement as a blind deconvolution problem with a convolution kernel of the form  $\exp(-\nu\|\xi\|^2)$ . If  $\nu$  were updated using the FSC at each iteration, this step would be equivalent to B-factor correction, where Wilson statistics first found applications in cryo-EM. This may suggest that B-factor correction should be performed within each iteration rather than post-processing.

## 5 Conclusion

We presented a prior distribution based on a simple yet expressive generative model for molecules. Then, we gave strategies to evaluate its hyperparameters and showed that, on simple inverse problems, the prior outperforms a prior similar to one often used in practice. Finally, we discussed properties of the Wilson prior, its connections to other regularization strategies, and stated potential implications for structure determination in cryo-EM.

## 6 Appendix

### 6.1 The Wilson prior in the spatial domain

We establish the representation of the covariance operator  $\Sigma_{\text{Wilson}}$  in the spatial domain:

$$(\Sigma_{\text{Wilson}}\phi)(x) = N_{\text{atoms}}f(x) * \left[ g(x) \left( \overline{f(x)} * \phi(x) \right) \right] - N_{\text{atoms}}f(x) * g(x) \left\langle g(x), \overline{f(x)} * \phi(x) \right\rangle \quad (16)$$

where  $*$  denotes convolution and  $\langle \cdot, \cdot \rangle$  is the standard inner product on  $L_2(\mathbb{R}^3)$ . We use the representation in the spectral domain established in [18],

$$(\hat{\Sigma}_{\text{Wilson}}\hat{\phi}_2)(\xi_1) = \int_{\xi_2} \hat{\Sigma}_{\text{Wilson}}(\xi_1, \xi_2) \hat{\phi}_2(\xi_2) d\xi_2 = \int_{\xi_2} N_{\text{atoms}} \hat{f}(\xi_1) \overline{\hat{f}(\xi_2)} \left( \hat{g}(\xi_1 - \xi_2) - \hat{g}(\xi_1) \overline{\hat{g}(\xi_2)} \right) \hat{\phi}_2(\xi_2) d\xi_2. \quad (17)$$

Let  $\phi_1(x), \phi_2(x) \in L^2(\mathbb{R}^3)$ , and  $\mathcal{F}$  denote the three-dimensional Fourier transform, then:

$$\begin{aligned} \langle \phi_1(x), \Sigma_{\text{Wilson}}\phi_2(x) \rangle &= \langle \mathcal{F}^{-1}\mathcal{F}\phi_1(x), \Sigma_{\text{Wilson}}\mathcal{F}^{-1}\mathcal{F}\phi_2(x) \rangle \\ &= \langle \mathcal{F}\phi_1(x), (\mathcal{F}\Sigma_{\text{Wilson}}\mathcal{F}^{-1})\mathcal{F}\phi_2(x) \rangle \end{aligned} \quad (18)$$

$$= \langle \hat{\phi}_1(\xi_1), \hat{\Sigma}_{\text{Wilson}}\hat{\phi}_2(\xi_1) \rangle \quad (19)$$

$$= N_{\text{atoms}} \left\langle \overline{\hat{f}(\xi_1)} \hat{\phi}_1(\xi_1), \int_{\xi_2} \overline{\hat{f}(\xi_2)} \left( \hat{g}(\xi_1 - \xi_2) - \hat{g}(\xi_1) \overline{\hat{g}(\xi_2)} \right) \hat{\phi}_2(\xi_2) d\xi_2 \right\rangle \quad (20)$$

$$\begin{aligned} &= N_{\text{atoms}} \left\langle \overline{\hat{f}(\xi_1)} \hat{\phi}_1(\xi_1), \int_{\xi_2} \hat{g}(\xi_1 - \xi_2) \overline{\hat{f}(\xi_2)} \hat{\phi}_2(\xi_2) d\xi_2 \right\rangle \\ &\quad - N_{\text{atoms}} \left\langle \overline{\hat{f}(\xi_1)} \hat{\phi}_1(\xi_1), \hat{g}(\xi_1) \right\rangle \left\langle \hat{g}(\xi_2), \overline{\hat{f}(\xi_2)} \hat{\phi}_2(\xi_2) \right\rangle \\ &= N_{\text{atoms}} \left\langle \phi_1(x), f(x) * \left( g(x) \left( \overline{f(x)} * \phi_2(x) \right) \right) \right\rangle \\ &\quad - N_{\text{atoms}} \langle \phi_1(x), f(x) * g(x) \rangle \left\langle g(x), \overline{f(x)} * \phi_2(x) \right\rangle \end{aligned} \quad (21)$$

Equation (18) follows from the unitary property of the Fourier transform  $\mathcal{F}^* = \mathcal{F}^{-1}$ , eq. (19) follows from the identity  $\text{Cov}(A\phi) = A\text{Cov}(\phi)A^*$ , eq. (20) follows from eq. (17), and eq. (21) follows from the convolution theorem. Since  $\phi_1, \phi_2 \in L_2(\mathbb{R}^3)$  are arbitrary, this establishes eq. (16). In quasi-matrix notation, the covariance operator is written as:

$$\Sigma_{\text{Wilson}} = N_{\text{atoms}} C_f (D_g - g g^*) C_f^*$$

where  $C_f$  denotes the convolution operator with  $f$ ,  $D_g$  is the multiplication operator with  $g$ , and the outer product is defined as  $(uv^*)\phi = u \langle v, \phi \rangle$ . In the denoising problem ( $A = I$ ) and with the choice of  $f(x) = \delta(x)$ , which implies  $C_f = I$ , the Wiener filter matrix in eq. (10) expressed in real space is:

$$H_{\text{Wilson}} = N_{\text{atoms}}(D_g - gg^*)(\Sigma_{\text{Wilson}} + \sigma^2 I)^{-1}$$

Note that in the case,  $g(x)$  acts as a mask for the MAP-Wilson estimate:

$$\phi_{\text{Wilson}}^{\text{MAP}} = \mu_{\text{Wilson}} + H_{\text{Wilson}}(y - \mu_{\text{Wilson}}) = g \cdot w, \text{ where } w = N_{\text{atoms}}\mathbb{1} + N_{\text{atoms}}(I - \mathbb{1}g^*)(\Sigma_{\text{Wilson}} + \sigma^2 I)^{-1}(y - \mu_{\text{Wilson}}),$$

where  $\mathbb{1}(x) = 1$ . That is, the MAP-Wilson estimate factorizes into to some function  $w$  multiplied by the shape function  $g$ , which implies the support of  $g$  contains the support of  $\phi_{\text{Wilson}}^{\text{MAP}}$ .

## 6.2 The diagonal prior as the uniform limit of the Wilson prior

Recall that the Wilson and diagonal prior are defined as follows:

$$\begin{aligned} \mu_{\text{diagonal}}(\xi) &= 0, \\ \Sigma_{\text{diagonal}}(\xi_1, \xi_2) &= \begin{cases} \text{PS}(\|\xi_1\|) & \text{if } \xi_1 = \xi_2, \\ 0 & \text{else} \end{cases}, \\ \hat{\mu}_{\text{Wilson}}(\xi) &= N_{\text{atoms}}\hat{f}(\xi)\hat{g}(\xi), \\ \hat{\Sigma}_{\text{Wilson}}(\xi_1, \xi_2) &= N_{\text{atoms}}\hat{f}(\xi_1)\overline{\hat{f}(\xi_2)}\left(\hat{g}(\xi_1 - \xi_2) - \hat{g}(\xi_1)\overline{\hat{g}(\xi_2)}\right). \end{aligned}$$

The diagonal of the diagonal covariance prior is the spherically-averaged power spectrum of the molecule. Under the bag of atoms model, the expected power spectrum is (see [18] for derivation):

$$\begin{aligned} \text{PS}(k) &= E \left[ \frac{1}{4\pi} \int_{S^2} |\hat{\phi}(k\omega)|^2 d\omega \right] \\ &= |\hat{f}(k)|^2 \left( N_{\text{atoms}} + N_{\text{atoms}}(N_{\text{atoms}} - 1) \frac{1}{4\pi} \int_{S^2} |\hat{g}(k\omega)|^2 d\omega \right) \end{aligned}$$

where we have assumed that  $f(x)$  is a spherically symmetric function. We now make the dependence on  $\hat{g}$  explicit with superscripts, and even though that choice is not properly defined, we proceed with  $\hat{g} = \delta$ . In that case, the mean and covariances of Wilson and diagonal agree for  $\xi \neq 0$ :

$$\hat{\mu}_{\text{Wilson}}^{\hat{g}=\delta}(\xi) = N_{\text{atoms}}\hat{f}(\xi)\delta(\xi) = 0 = \mu_{\text{diagonal}}^{\hat{g}=\delta}(\xi),$$

and similarly for  $\xi_1, \xi_2 \neq 0$ :

$$\begin{aligned} \hat{\Sigma}_{\text{Wilson}}^{\hat{g}=\delta}(\xi_1, \xi_2) &= \begin{cases} N_{\text{atoms}}|\hat{f}(\|\xi_1\|)|^2 & \text{if } \xi_1 = \xi_2, \\ 0 & \text{else} \end{cases}, \\ \hat{\Sigma}_{\text{diagonal}}^{\hat{g}=\delta}(\xi_1, \xi_1) &= |\hat{f}(\|\xi_1\|)|^2 \left( N_{\text{atoms}} + N_{\text{atoms}}(N_{\text{atoms}} - 1) \frac{1}{4\pi} \int_{S^2} |\delta(k\omega)|^2 d\omega \right) = N_{\text{atoms}}|\hat{f}(\|\xi_1\|)|^2, \end{aligned}$$

thus  $\hat{\Sigma}_{\text{Wilson}}^{\hat{g}=\delta}(\xi_1, \xi_2) = \hat{\Sigma}_{\text{diagonal}}^{\hat{g}=\delta}(\xi_1, \xi_2)$ .  $\delta$  is not a valid choice of  $\hat{g}$  as there is no well-defined probability density function  $g$  for which  $\mathcal{F}(g) = \delta(x)$ . Instead, we consider uniform distribution on a ball of radius  $R$ :  $g_R(x) = \frac{3}{4\pi R^3} \mathcal{X}_B(x/R)$  where  $\mathcal{X}_B$  is the characteristic function of the unit ball, and let the radius grow to infinity. The Fourier transform of  $g_R$  is:

$$\hat{g}_R(\xi) = -\frac{3 \cos(2\pi R|\xi|)}{4\pi^2 R^2 |\xi|^2} + \frac{3 \sin(2\pi R|\xi|)}{8\pi^3 R^3 |\xi|^3}$$

and we have the desired limit property  $\lim_{R \rightarrow \infty} \hat{g}_R(x) = \delta(x)$ . Consequently:

$$\begin{aligned} \lim_{R \rightarrow \infty} \hat{\mu}_{\text{Wilson}}^{\hat{g}=\hat{g}_R}(\xi) &= \lim_{R \rightarrow \infty} \mu_{\text{diagonal}}^{\hat{g}=\hat{g}_R}(\xi) \quad \text{for } \xi \neq 0, \\ \lim_{R \rightarrow \infty} \hat{\Sigma}_{\text{Wilson}}^{\hat{g}=\hat{g}_R}(\xi_1, \xi_2) &= \lim_{R \rightarrow \infty} \Sigma_{\text{diagonal}}^{\hat{g}=\hat{g}_R}(\xi_1, \xi_2) \quad \text{for } \xi_1, \xi_2 \neq 0. \end{aligned}$$

That is, the diagonal and Wilson priors agree in the limit of a uniform distribution on the entire domain at non-zero frequencies.

## Acknowledgement

M.A.G. and A.S. are supported in part by AFOSR FA9550-20-1-0266, the Simons Foundation Math+X Investigator Award, NSF BIGDATA Award IIS-1837992, NSF DMS-2009753, and NIH/NIGMS 1R01GM136780-01. We thank Eric J. Verbeke for valuable discussion and help in generating figures.

## References

- [1] Barnett, A.H.: Aliasing error of the  $\exp(\beta\sqrt{1-z^2})$  kernel in the nonuniform fast fourier transform. *Applied and Computational Harmonic Analysis* **51**, 1–16 (2021)
- [2] Barnett, A.H., Magland, J., af Klinteberg, L.: A parallel nonuniform fast Fourier transform library based on an “exponential of semicircle” kernel. *SIAM Journal on Scientific Computing* **41**(5), C479–C504 (2019)
- [3] Bendory, T., Bartesaghi, A., Singer, A.: Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE signal processing magazine* **37**(2), 58–76 (2020)
- [4] Bertsekas, D.P.: *Constrained optimization and Lagrange multiplier methods*. Academic press (2014)
- [5] Buzzard, G.T., Chan, S.H., Sreehari, S., Bouman, C.A.: Plug-and-play unplugged: Optimization-free reconstruction using consensus equilibrium. *SIAM Journal on Imaging Sciences* **11**(3), 2001–2020 (2018)
- [6] Gramacki, A.: *Nonparametric kernel density estimation and its computational aspects*. Springer (2018)
- [7] Hestenes, M.R., Stiefel, E., et al.: *Methods of conjugate gradients for solving linear systems*, vol. 49. NBS Washington, DC (1952)
- [8] Kimanius, D., Zickert, G., Nakane, T., Adler, J., Lunz, S., Schönlieb, C.B., Öktem, O., Scheres, S.H.: Exploiting prior knowledge about biological macromolecules in cryo-EM structure determination. *IUCrJ* **8**(1), 60–75 (2021)
- [9] Peng, L.M., Ren, G., Dudarev, S., Whelan, M.: Robust parameterization of elastic and absorptive electron atomic scattering factors. *Acta Crystallographica Section A: Foundations of Crystallography* **52**(2), 257–276 (1996)
- [10] Punjani, A., Brubaker, M.A., Fleet, D.J.: Building proteins in a day: Efficient 3D molecular structure estimation with electron cryomicroscopy. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(4), 706–718 (2016)
- [11] Punjani, A., Rubinstein, J.L., Fleet, D.J., Brubaker, M.A.: cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination. *Nature methods* **14**(3), 290–296 (2017)
- [12] Punjani, A., Zhang, H., Fleet, D.J.: Non-uniform refinement: adaptive regularization improves single-particle cryo-EM reconstruction. *Nature methods* **17**(12), 1214–1221 (2020)
- [13] Reehorst, E.T., Schniter, P.: Regularization by denoising: Clarifications and new interpretations. *IEEE transactions on computational imaging* **5**(1), 52–67 (2018)
- [14] Romano, Y., Elad, M., Milanfar, P.: The little engine that could: Regularization by denoising (RED). *SIAM Journal on Imaging Sciences* **10**(4), 1804–1844 (2017)
- [15] Rosenthal, P.B., Henderson, R.: Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *Journal of molecular biology* **333**(4), 721–745 (2003)
- [16] Scheres, S.H.: A Bayesian view on cryo-EM structure determination. *Journal of molecular biology* **415**(2), 406–418 (2012)
- [17] Scheres, S.H.: RELION: implementation of a Bayesian approach to cryo-EM structure determination. *Journal of structural biology* **180**(3), 519–530 (2012)
- [18] Singer, A.: Wilson statistics: derivation, generalization and applications to electron cryomicroscopy. *Acta crystallographica. Section A, Foundations and advances* **77**(Pt 5), 472–479 (2021). DOI 10.1107/s205327332100752x. URL <https://doi.org/10.1107/S205327332100752X>
- [19] Singer, A., Sigworth, F.J.: Computational methods for single-particle electron cryomicroscopy. *Annual Review of Biomedical Data Science* **3**, 163–190 (2020)
- [20] Venkatakrishnan, S.V., Bouman, C.A., Wohlberg, B.: Plug-and-play priors for model based reconstruction. In: 2013 IEEE Global Conference on Signal and Information Processing, pp. 945–948. IEEE (2013)

- [21] Wang, W., Carreira-Perpinán, M.A.: Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application. arXiv preprint arXiv:1309.1541 (2013)
- [22] Warshamanage, R., Yamashita, K., Murshudov, G.N.: Emda: A python package for electron microscopy data analysis. *Journal of Structural Biology* **214**, 107,826 (2021)
- [23] Wilson, A.J.C.: Determination of absolute from relative X-ray intensity data. *Nature* **150**(3796), 151–152 (1942)