

On some properties of the off-grid Normalized Matched Filter

S. Trottier^{*†}, J. Bosse^{*}, O. Rabaste^{*}, P. Forster[‡], J.-P. Ovarlez^{*†}

^{*}DEMR, ONERA, Université Paris-Saclay, F-91120 Palaiseau, France.

[†]SONDRA CentraleSupélec, Université Paris-Saclay, F-91192 Gif-sur-Yvette, France.

[‡]Université Paris-Saclay, ENS Paris-Saclay, CNRS, SATIE, F-91190, Gif-sur-Yvette, France.

Abstract—Detecting a target signal in noise is a central problem in radar signal processing. The Normalized Matched Filter (NMF) is widely used, typically evaluated on a discrete grid of target parameters. In practice, targets rarely align with grid points, reducing detection performance. The off-grid NMF addresses this by maximizing over all possible parameters, improving sensitivity to off-grid targets. Building on recent results establishing the relationship between the threshold of the off-grid test and the false alarm probability, we investigate properties of the off-grid NMF that does not seem to be well known in the literature and illustrate through simulations that it exhibits an almost constant false alarm rate in realistic radar scenarios.

Index Terms—Radar, Detection, Off-grid, Normalized Matched Filter

I. INTRODUCTION

A fundamental problem in signal processing, encountered in many areas such as radar, telecommunications, and seismology, is to decide whether an observed signal contains a target signal of interest, hypothesis (\mathcal{H}_1) or not, hypothesis (\mathcal{H}_0). Detection theory provides several frameworks to address this problem, among which the likelihood ratio test (LRT) remains the reference approach [1]. However, the LRT requires full knowledge of all scene parameters, such as the noise covariance matrix or the target steering vector, which is rarely achievable in practice. To address this limitation, the Generalized Likelihood Ratio Test (GLRT) was introduced, replacing unknown parameters with their maximum likelihood estimates [2]. This principle has been widely adopted in radar detection, where environmental parameters often vary with time, space, and frequency.

Among the various detectors derived from the GLRT, the Normalized Matched Filter (NMF) [3], also known as the adaptive cosine estimator [4], is one of the most popular. It provides a normalized measure of similarity between the received signal and a hypothesized target steering vector. Because target parameters such as Doppler frequency or angle of arrival are generally unknown, the detection process typically involves evaluating the NMF over a discrete grid of possible parameter values. The resulting implementation is widely used in practice probably due to its low computational cost and compatibility with real-time radar systems, where both time and processing resources are constrained. Consequently, most studies in the literature describe the statistical properties of the NMF for a single grid point, assuming that the target is aligned with one of the grid hypotheses.

In real-world scenarios, however, target parameters have no reason to coincide with the grid points used for the detection tests. This mismatch can lead to significant performance degradation, as the NMF becomes less sensitive to off-grid targets [5], [6], [7]. A natural way to mitigate this issue is to consider the maximum of the NMF over all possible target parameters instead of restricting the search to discrete grid points, leading to what we call the off-grid NMF. While such a continuous optimization problem might appear computationally demanding, some efficient solutions exist [8]. In addition, the statistical properties of the off-grid detector—particularly the relationship between the false alarm probability (P_{fa}) and the detection threshold—have recently been established, [9]. Interestingly, while the conventional NMF is said to have a constant false alarm rate (CFAR) property, this does not hold in the off-grid setting. It is only true for a single test point.

In this paper, we will first show the drawbacks of the NMF detector when applied to a sampled grid of the parameter space: in particular, the tests corresponding to the different sampling points are not independent, thus leading to difficulties in computing the false alarm probability and to a non-CFAR test. This is why we then turn our attention to the off-grid NMF: we provide an in-depth study of this off-grid test, emphasizing how the P_{fa} -threshold relationship can be derived analytically [10], [11]. We provide further details on this derivation, showing that it corresponds to the expected number of false alarms. This allows us to derive several important properties of the off-grid NMF, particularly in multi-target scenarios. Finally, simulations show that the off-grid NMF exhibits an almost CFAR behavior in practice.

The paper is structured as follows. Section II presents the problem formulation. In Section III, we provide precision on the derivation of the P_{fa} -threshold relationship for the off-grid NMF. Our main contributions are presented in Section IV, which discusses the implications of this relationship and presents numerical experiments illustrating the nearly CFAR behavior of the off-grid NMF.

Notations: *Italic type* indicates a scalar quantity, lower case **boldface** indicates a vector quantity, and upper case **boldface** indicates a matrix or tensor. The transpose and transpose conjugate operators are T and H respectively. $\mathbf{x} \sim \mathcal{CN}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ is a random complex circular Gaussian vector with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. The derivative of a single-

parameter function f is denoted \hat{f} .

II. PROBLEM FORMULATION

In radar detection, the task is to decide between two possible scenarios: the received signal, modeled as a complex vector \mathbf{y} of dimension N , is either pure noise \mathbf{n} or contains a target signal $c\mathbf{a}(\theta)$ embedded in noise \mathbf{n} . Formally, the detection problem can be expressed as:

$$\begin{cases} \mathcal{H}_0 : \mathbf{y} = \mathbf{n}, & (\text{noise only}) \\ \mathcal{H}_1 : \mathbf{y} = c\mathbf{a}(\theta) + \mathbf{n}, & (\text{signal plus noise}) \end{cases} \quad (1)$$

where the noise vector is assumed to be circular complex Gaussian, $\mathbf{n} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{R})$. Here, \mathbf{R} is known while the noise power σ^2 is unknown [12]. The complex scalar c represents the unknown target amplitude, and $\mathbf{a}(\theta)$ is a steering vector parameterized by an unknown Doppler frequency θ , defined as:

$$\mathbf{a}(\theta) = \frac{1}{\sqrt{N}} (1, e^{2i\pi\theta}, \dots, e^{2i\pi(N-1)\theta})^T, \quad \theta \in \Theta = [0, 1]. \quad (2)$$

Under the null hypothesis \mathcal{H}_0 , the distribution of \mathbf{y} depends solely on the parameter set $\lambda_0 = \{\sigma^2\}$, whereas under \mathcal{H}_1 , it depends on $\lambda_1 = \{c, \sigma^2, \theta\}$.

The generalized likelihood ratio test (GLRT) evaluates the ratio of the maximum likelihoods under the two hypotheses:

$$\frac{\max_{\lambda_1} f_{\mathcal{H}_1}(\mathbf{y})}{\max_{\lambda_0} f_{\mathcal{H}_0}(\mathbf{y})} \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \omega^2, \quad (3)$$

which leads to the well-known test:

$$\Lambda_{\text{NMF-off-grid}} = \max_{\theta \in \Theta} |\mathbf{s}^H(\theta) \mathbf{u}|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\geq}} \omega^2, \quad (4)$$

where $\mathbf{s}(\theta) = \frac{\mathbf{R}^{-1/2} \mathbf{a}(\theta)}{\|\mathbf{R}^{-1/2} \mathbf{a}(\theta)\|}$ and $\mathbf{u} = \frac{\mathbf{R}^{-1/2} \mathbf{y}}{\|\mathbf{R}^{-1/2} \mathbf{y}\|}$.

Because the off-grid NMF involves a maximization over θ , obtaining its distribution under \mathcal{H}_0 is challenging.

Therefore, only an approximation of this test is usually performed: the on-grid NMF that consists in computing the single value NMF over a predefined sampled set $\{\theta_1, \dots, \theta_N\} \in \Theta$. First, note that testing only one single value of θ leads to the well-known Normalized Matched Filter (NMF):

$$\Lambda_{\text{NMF}}(\theta) = |\mathbf{s}^H(\theta) \mathbf{u}|^2, \quad (5)$$

whose P_{fa} -threshold relationship is given by:

$$P_{\text{fa-NMF}} = (1 - \omega^2)^{N-1}. \quad (6)$$

The NMF (5) exhibits a constant false alarm rate in the sense that the P_{fa} -threshold relationship is independent of \mathbf{R} and the testing point θ . However, while this property holds for a single testing point, it no longer applies when the test is performed over the entire grid:

$$\Lambda_{\text{NMF-on-grid}} = \max_{\theta \in \{\theta_1, \dots, \theta_N\}} \Lambda_{\text{NMF}}(\theta), \quad (7)$$

where $\{\theta_1, \dots, \theta_N\}$ are the grid points. Concerning the probability of false alarm of the on-grid test, it is generally hard

to derive since the correlations between the tests on the grid points are not known in the general case.

Moreover, the NMF is selective since it corresponds to the measure of an angle between the received signal and the steering vector, and the on-grid test conserves this selectiveness around each grid point. This property is mostly depicted as positive in the literature since it allows this detector to be robust to high-power noise and interferences, but it has a major drawback: off-grid targets. When facing off-grid targets, the on-grid test performance is highly degraded, even with a signal-to-noise ratio going to infinity [6]. Due to this major drawback, in the following parts, we will focus the rest of our study on the off-grid NMF, which by construction is selective but also robust to off-grid targets.

III. OFF-GRID FALSE ALARM PROBABILITY

In [9], using a geometric approach inspired by [10], the authors derived an explicit expression for the false alarm probability

$$P_{\text{fa}} = \mathbb{P}(\Lambda_{\text{NMF-off-grid}} \geq \omega^2 \mid \mathcal{H}_0), \quad (8)$$

as a function of the threshold ω . In this section, we go into the details of the derivation of the P_{fa} -threshold relationship. We provide additional explanations beyond those presented in the original paper [9], establishing a direct link with the expected number of false alarms and the inherent detection limitations of the NMF in scenarios with multiple targets. We also discuss how the noise covariance matrix may impact this relationship.

First, in the same way as in [9], we note that:

$$|\mathbf{s}^H(\theta) \mathbf{u}|^2 = \max_{\alpha} \Re(e^{j\alpha} \mathbf{s}^H(\theta) \mathbf{u})^2, \quad (9)$$

and deduce:

$$\max_{\theta \in \Theta} |\mathbf{s}^H(\theta) \mathbf{u}|^2 = \left(\max_{\theta \in \Theta, \alpha} \gamma^T(\theta, \alpha) \tilde{\mathbf{u}} \right)^2, \quad (10)$$

with

$$\gamma(\theta, \alpha) = \cos(\alpha) \gamma_1(\theta) + \sin(\alpha) \gamma_2(\theta), \quad (11)$$

where

$$\gamma_1(\theta) = \begin{bmatrix} \Re(\mathbf{s}(\theta)) \\ \Im(\mathbf{s}(\theta)) \end{bmatrix}, \quad \gamma_2(\theta) = \begin{bmatrix} -\Im(\mathbf{s}(\theta)) \\ \Re(\mathbf{s}(\theta)) \end{bmatrix}, \quad (12)$$

and

$$\tilde{\mathbf{u}} = \begin{bmatrix} \Re(\mathbf{u}) \\ \Im(\mathbf{u}) \end{bmatrix}. \quad (13)$$

In the following, \mathbf{y} is noise only since we focus on the false alarm probability. In that case, $\mathbf{R}^{-1/2} \mathbf{y} \sim \mathcal{CN}(\mathbf{0}, \mathbf{I})$ and so $\tilde{\mathbf{u}}$ is uniformly distributed over the $2N$ -sphere \mathbb{S}^{2N-1} where N is the dimension of \mathbf{y} . Moreover, for $\theta \in \Theta, \alpha \in [0, 2\pi)$, $\gamma(\theta, \alpha)$ describes a certain manifold over the $2N$ -sphere:

$$\mathcal{M} = \{\gamma(\theta, \alpha); \theta \in \Theta, \alpha \in [0, 2\pi)\} \quad (14)$$

For a single parameter θ , \mathcal{M} is thus a 2D-manifold. Since the statistic of interest is the following:

$$\mathbb{P} \left(\max_{\theta \in \Theta, \alpha} \gamma^T(\theta, \alpha) \tilde{\mathbf{u}} \geq \omega \right), \quad (15)$$

we can see that there is a false alarm as soon as $\tilde{\mathbf{u}}$ is in the tube of radius $r = \sqrt{1 - \omega^2}$ around the manifold \mathcal{M} : $\mathcal{T}(\mathcal{M}, r)$. The proper definition of a tube involves the cross-sections that we will introduce later [11], but in our case, since the manifold \mathcal{M} is closed, we can write:

$$\mathcal{T}(\mathcal{M}, r) = \{ \mathbf{v} \in \mathbb{S}^{2N-1}; \text{distance}(\mathbf{v}, \mathcal{M}) \leq r \}. \quad (16)$$

Since $\tilde{\mathbf{u}}$ is uniformly distributed over the $2N$ -sphere, the ratio between the volume of $\mathcal{T}(\mathcal{M}, r)$ and the volume of \mathbb{S}^{2N-1} is exactly the P_{fa} . Thus, we will now focus on computing the volume of $\mathcal{T}(\mathcal{M}, r)$.

In [10], Weyl proposes a general framework to compute the volume of a tube around a manifold. The volume of $\mathcal{T}(\mathcal{M}, r)$ is computed using the cross-sections of the tube. A cross-section associated to a point $\gamma(\theta, \alpha)$ of the manifold is:

$$\mathcal{C}(\gamma(\theta, \alpha)) = \left\{ \mathbf{v} \in \mathcal{T}(\mathcal{M}, r); \mathbf{v} - \gamma(\theta, \alpha) \perp \frac{\partial \gamma(\theta, \alpha)}{\partial \theta}, \right. \\ \left. \mathbf{v} - \gamma(\theta, \alpha) \perp \frac{\partial \gamma(\theta, \alpha)}{\partial \alpha} \right\}. \quad (17)$$

Figure 1 represents a tube around a manifold, and a cross-section associated with a certain point is also illustrated. As one can see, a tube is basically the union of all cross-sections. The formula provided by Weyl states that the volume of $\mathcal{T}(\mathcal{M}, r)$ can be computed as the sum of the volumes of the cross-sections. However, if the radius r of the tube is above a certain level, some cross-sections may intersect. In this case, the total volume of the tube computed using Weyl's formula becomes overestimated.

We give some details about this overestimation. Let us consider that there is a detected target with parameters (θ, α) when $\gamma^T(\theta, \alpha)\tilde{\mathbf{u}}$ corresponds to a local maximum above the threshold ω . Thus, we notice that when $\tilde{\mathbf{u}}$ belongs to multiple cross-sections, for instance $\mathcal{C}(\gamma(\theta_1, \alpha_1))$ and $\mathcal{C}(\gamma(\theta_2, \alpha_2))$, this means that there are multiple detected targets. Hence, computing the volume of all cross-sections does not lead to compute the probability of false alarm but it can be seen that it corresponds rather to the expectation of the number of false alarms. However, asymptotically, i.e., for a radius r small enough corresponding to a threshold ω high enough, the cross sections do not overlap and Weyl's expression becomes exact.

The P_{fa} -threshold relationship for the off-grid NMF, which is asymptotically exact as ω grows, has been proposed in [9] and is expressed as:

$$P_{\text{fa}} = \frac{\Gamma(N)\omega(1 - \omega^2)^{N - \frac{3}{2}}}{\pi^{\frac{1}{2}}\Gamma(N - \frac{1}{2})} \int_0^1 \sqrt{\|\dot{\gamma}_1\|^2 - (\dot{\gamma}_1^T \gamma_2)^2} d\theta, \quad (18)$$

where Γ corresponds to the gamma function and γ_1 and γ_2 are defined in (12). Since they both depend on \mathbf{R} , this expression depends on the noise covariance matrix. It can be simplified under white noise to obtain [9]:

$$P_{\text{fa}} = \sqrt{\frac{\pi(N^2 - 1)}{3}} \frac{\Gamma(N)\omega(1 - \omega^2)^{N - 3/2}}{\Gamma(N - 1/2)}. \quad (19)$$




Figure 1. Illustration of a blue tube of radius r around a closed 1D-manifold \mathcal{M} in black with a cross-section in red.

If we look closely into the P_{fa} -threshold relationship in equation (18), we can split it into three terms:

$$\begin{cases} \text{Volume of a cross-section} : \frac{\pi^{N-3/2}}{\Gamma(N-1/2)}\omega(1 - \omega^2)^{N-3/2}, \\ \text{Volume of } \mathbb{S}^{2N-1} : \frac{2\pi^N}{\Gamma(N)}, \\ \text{"Length" of } \mathcal{M} : 2\pi \int_0^1 \sqrt{\|\dot{\gamma}_1\|^2 - (\dot{\gamma}_1^T \gamma_2)^2} d\theta. \end{cases}$$

Finally, it is clear that the volume of the tube corresponds to the volume of a cross-section times the length of the manifold, and to obtain the P_{fa} , one must divide the expression by the volume of \mathbb{S}^{2N-1} . For the remainder of the article, we will denote the "length" of the manifold by:

$$L(\mathbf{R}) = 2\pi \int_0^1 \sqrt{\|\dot{\gamma}_1\|^2 - (\dot{\gamma}_1^T \gamma_2)^2} d\theta.$$

We present the analytical P_{fa} -threshold relationship under white noise conditions, given in (19), and compare it with Monte-Carlo simulations respectively computing the expectation of the number of false alarms and the probability of false alarm in Figure 2 for $N = 16$. For each threshold value, the number of Monte-Carlo trials is set to $\frac{400}{\text{Theoretical } P_{\text{fa}}}$, with a minimum of 10000. The results show that, for false alarm probabilities below 10^{-1} , the analytical expression closely matches the simulation outcomes. This is because the formula is only asymptotically exact, i.e., when the tube does not self-overlap. More generally, the analytical expression coincides with the expected number of false alarms, which asymptotically is equal to the P_{fa} .

IV. PROPERTIES OF THE OFF-GRID NMF

In this section, we discuss several properties of the off-grid NMF—some directly derived from the preceding analysis—that, to the best of our knowledge, have not been explicitly reported in the existing literature.

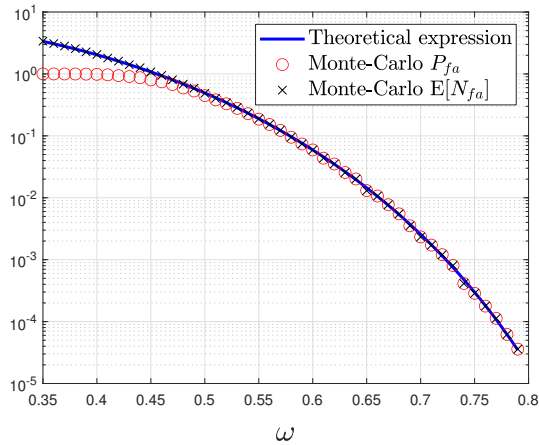


Figure 2. Comparison between the theoretical P_{fa} -threshold relationship (19) and Monte-Carlo simulations of the expected number of false alarms $E[N_{fa}]$ and probability of false alarm for $N = 16$, $\mathbf{R} = \mathbf{I}$.

A. Only one target can be detected over Θ !

When analyzing the detection zone associated with the NMF, i.e., a tubular neighborhood around \mathcal{M} on the $2N$ -sphere, one can observe that for a sufficiently high threshold ω , the tube's cross-sections do not overlap. In other words, each point within the tube corresponds uniquely to a single cross-section. Consequently, for any vector $\tilde{\mathbf{u}}$, the function $f(\theta, \alpha) = \gamma^T(\theta, \alpha)\tilde{\mathbf{u}}$ admits at most one maximum exceeding the threshold ω . Therefore, no more than one target can be detected. This property is inherent to the NMF in general and is not exclusive to its off-grid formulation.

For instance, if we look at the white noise setting for the on-grid NMF, with a grid spacing matching the Doppler resolution $\frac{1}{N}$, we know that $\mathbf{s}(\theta_1), \dots, \mathbf{s}(\theta_N)$ is an orthonormal basis of \mathbb{C}^N . Let $\mathbf{u} = [u_1, \dots, u_N]^T \in \mathbb{C}^N$ such that \mathbf{u} is uniformly distributed over the complex N -sphere \mathbb{S}_C^{N-1} . A classical result in statistics is:

$$\left(|u_1|^2, \dots, |u_N|^2\right) \sim Dir(1, 1, \dots, 1), \quad (20)$$

where $Dir(\mathbf{x})$ corresponds to a Dirichlet distribution of parameter \mathbf{x} . Replacing the canonical orthonormal basis by $\mathbf{s}(\theta_1), \dots, \mathbf{s}(\theta_N)$, we deduce:

$$\left(|\mathbf{s}^H(\theta_1)\mathbf{u}|^2, \dots, |\mathbf{s}^H(\theta_N)\mathbf{u}|^2\right) \sim Dir(1, 1, \dots, 1), \quad (21)$$

Hence, the different tests $\Lambda_{\text{NMF}}(\theta_1), \dots, \Lambda_{\text{NMF}}(\theta_N)$ are negatively correlated by the property of the Dirichlet distribution and, therefore, are not independent.

Another noteworthy property is the threshold limit $\omega(k)$, above which no more than k targets can be detected in this context. This corresponds to distributing the power of \mathbf{u} uniformly across $k+1$ vectors of the orthonormal basis, yielding $\omega(k) = \frac{1}{\sqrt{k+1}}$. Hence, if the threshold is above $\frac{1}{\sqrt{2}}$, at most one target can be detected by the on-grid NMF (under white noise).

This property may appear counterintuitive, as we are generally more accustomed to the conventional Matched Filter (MF), which can, in principle, detect up to N distinct targets—one per resolution cell—provided each target has sufficiently high power. In typical radar processing, the Doppler frequency domain $\Theta = [0, 1)$ is discretized into resolution cells, $\left\{ \left[0, \frac{1}{N}\right), \left[\frac{1}{N}, \frac{2}{N}\right), \dots, \left[\frac{N-1}{N}, 1\right) \right\}$, and each cell is associated with either a detection or a non-detection.

In contrast, for the NMF, in low-dimensional scenarios commonly encountered in radar detection ($N \leq 20$) and under low false-alarm probability regimes ($P_{fa} \leq 10^{-4}$), at most one target can be reliably detected over the entire parameter space Θ . This is because, in these scenarios, the threshold ω is above $\frac{1}{\sqrt{2}}$. Consequently, partitioning Θ into independent resolution cells is unnecessary, as these cells are not statistically independent in the NMF framework. Instead, the detection test can be performed directly over the continuous domain Θ . An efficient implementation of this off-grid detection procedure is provided in [8].

B. Impact of multiple targets

For the conventional Matched Filter, when multiple targets are present and sufficiently strong and well separated, all targets are likely to be detected. In the worst-case scenario, at least the target with the highest SNR will be detected, and its detection performance is generally unaffected by the presence of other targets, provided that they are adequately separated. In contrast, for the normalized matched filter (NMF), this property no longer holds. This can be explained by the previously noted fact that, above a certain threshold, the NMF can detect at most one target.

Even when targets are well separated, the normalization introduces an effect analogous to mutual interference. Consequently, if only one target can be detected, the presence of additional targets reduces the probability of detecting the one with the highest SNR. This behavior is counterintuitive, as it does not occur with the conventional Matched Filter.

To illustrate this phenomenon, we performed simulations in a scenario with two targets, varying the ratio of their SNRs. With $N = 16$ and under white noise, we compare the target detection probability for a threshold ω corresponding to an analytical false alarm probability $P_{fa} = 10^{-4}$ in a single target context (at frequency $f_1 = 0.5$) with the case where another target is present at a lower SNR and at a different frequency ($f_2 = 0.8 \pm \mathbf{U}$ with $\mathbf{U} \sim \mathcal{U}(0, 0.1)$). In a white noise context, the SNR of a target is defined as:

$$\text{SNR} = \frac{N|c|^2}{\sigma^2}, \quad (22)$$

with c the complex amplitude of the target and σ^2 the variance of the noise. The relative SNR values considered between the main target and the other are 10 dB, 5 dB, 3 dB, and 0 dB. We consider that we detect the main target if the off-grid NMF is above the threshold ω and if the corresponding Doppler frequency is close enough to the reference f_1 , where

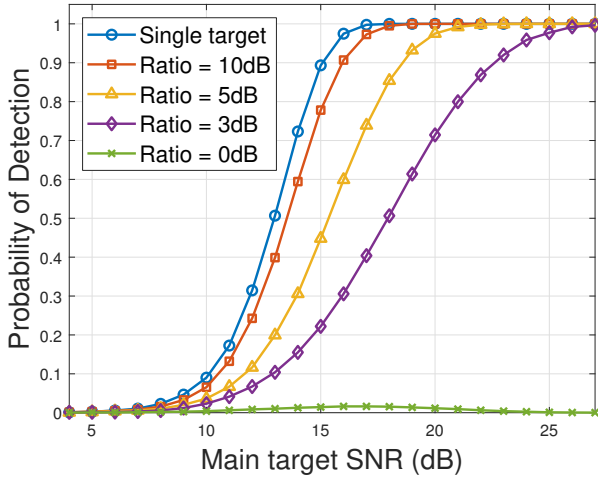


Figure 3. Detection probability of the main target with respect to its SNR in the presence of a second target with an SNR lower by a factor “Ratio” for different ratio values, for a threshold corresponding to $P_{fa} = 10^{-4}$.

close enough means at most $0.5/N$, so half a resolution cell. The results in Figure 3 clearly show that the presence of another significant target degrades the detection of the main one, making it undetectable when both targets have the same SNR. Note also that in all these scenarios, the secondary target is not detected. In particular, these effects are not specific to the off-grid NMF, as the on-grid NMF exhibits the same behavior.

C. Impact of the noise covariance matrix

In the P_{fa} -threshold relationship of equation (18) corresponding to the off-grid NMF, the covariance matrix \mathbf{R} appears only through the length of the manifold \mathcal{M} . Consequently, it can affect the relationship merely by a constant multiplicative factor, while the overall functional dependence between P_{fa} and the threshold remains essentially independent of \mathbf{R} .

Nevertheless, it is important to assess the possible variability of the manifold length. In the general case, when the only constraint on \mathbf{R} is positive definiteness, one can construct non-physical or “pathological” examples that yield arbitrarily small manifold lengths and quite large ones. For instance, let us consider the following diagonal matrix:

$$\mathbf{M}_1 = \text{diag}(1 + \epsilon, \epsilon, \epsilon, \dots, \epsilon). \quad (23)$$

If $\mathbf{R}^{-1/2} = \mathbf{M}_1$, then it can be shown that:

$$L(\mathbf{R}) = \frac{4\pi^2\epsilon \sqrt{\sum_{k=1}^{N-1} k^2 - \frac{\epsilon^2 \left(\sum_{k=1}^{N-1} k\right)^2}{1 + 2\epsilon + N\epsilon^2}}}{\sqrt{1 + 2\epsilon + N\epsilon^2}}.$$

Thus, by choosing ϵ sufficiently small, the manifold length can be made arbitrarily small.

Now, let us consider another diagonal matrix:

$$\mathbf{M}_2 = \text{diag}(1 + \epsilon, \epsilon, \dots, \epsilon, 1 + \epsilon). \quad (24)$$




Figure 4. Comparison of $l(\mathbf{R}(\rho), \theta)$ for different values of ρ .

If we now consider $\mathbf{R}^{-1/2} = \mathbf{M}_2$, then we get:

$$L(\mathbf{R}) = 2\pi^2(N - 1) + O(\epsilon). \quad (25)$$

This should be contrasted with the white-noise case, where we have:

$$L(\mathbf{I}) = 2\pi^2 \sqrt{\frac{N^2 - 1}{3}}. \quad (26)$$

Hence, $L(\mathbf{R})$ from (25) and $L(\mathbf{I})$ differ by nearly a factor of $\sqrt{3}$, which represents the maximum ratio achievable when the covariance matrix is restricted to be diagonal.

In practice, the noise covariance matrix captures temporal correlations between successive samples. Under the assumption of stationary noise, it can therefore be well approximated by a Toeplitz structure. Although an analytical proof that the total length of the manifold \mathcal{M} remains nearly invariant in this setting is not provided, the numerical results in Figures 5 and 6 show that the off-grid NMF exhibits an almost-CFAR behavior. As a result, the simplified P_{fa} -threshold relationship in (19) provides a reliable approximation, largely independent of the actual noise covariance matrix.

The primary difficulty in deriving an analytical proof arises from the strong dependence of $l(\mathbf{R}, \theta) = \sqrt{\|\hat{\gamma}_1\|^2 - (\hat{\gamma}_1^T \gamma_2)^2}$ on both the noise covariance matrix \mathbf{R} and the parameter θ . As a result, establishing an upper bound for the integral is challenging, since such a bound cannot be obtained directly by bounding $l(\mathbf{R}, \theta)$.

In Figure 4, we plot the value of $l(\mathbf{R}, \theta)$ with respect to θ for different covariance matrix $\mathbf{R}(\rho)$ such that:

$$\mathbf{R}(\rho) = \text{Toep}([1, \rho, \rho^2, \dots, \rho^{N-1}]), \quad (27)$$

where Toep denotes the Toeplitz matrix operator. This model is widely used in radar signal processing, particularly for representing clutter in stationary radar environments. As previously discussed and illustrated in Figure 4, the quantity $l(\mathbf{R}, \theta)$ may vary significantly with \mathbf{R} and θ . However, the integral of




Figure 5. Ratio $\frac{L(\mathbf{R}(\rho))}{L(\mathbf{I})}$ with respect to ρ .




Figure 6. Histogram of values of the ratio $\frac{L(\mathbf{R})}{L(\mathbf{I})}$ with \mathbf{R} corresponding to the 100000 generated Toeplitz matrices.

$l(\mathbf{R}, \theta)$ over θ exhibits only limited variation with respect to \mathbf{R} , as shown in Figure 5.

Taking $\rho \in [0, 1)$ and $N = 16$, the results presented in Figure 5 correspond to the ratio:

$$\frac{L(\mathbf{R}(\rho))}{L(\mathbf{I})}, \quad (28)$$

with respect to ρ . We observe that even in the extreme case where ρ approaches one, the increase in the manifold length remains below 2.5%.

We also generated 100,000 random covariance matrices \mathbf{R} , constructed as positive-definite complex Toeplitz matrices derived from the autocorrelation of a multi-peak Gaussian spectrum. This procedure yields a more general and realistic clutter Power Spectral Density model:

$$S(f) = \sum_i w_i \exp\left(-\frac{(f - f_i)^2}{2\sigma_i^2}\right). \quad (29)$$

In this case, each spectrum is generated by drawing five frequencies f_i independently and uniformly over $[0, 1]$. For

each spectral component, the width σ_i is sampled from a uniform distribution $\mathcal{U}(0.01, 0.1)$, and the amplitude w_i from $\mathcal{U}(0, 10)$.

Figure 6 shows a histogram of the values obtained for the ratio $\frac{L(\mathbf{R})}{L(\mathbf{I})}$, where each \mathbf{R} is generated as described previously. We notice that the length of the manifold rarely increases by more than 5%. Thus, the dependence in \mathbf{R} of the P_{fa} -threshold relationship is limited. Another property that could be of theoretical interest is that the minimum manifold length for noise covariance matrices that are Toeplitz appears to be achieved when $\mathbf{R} = \mathbf{I}$. There was no apparent reason to expect this, since, for instance, it does not hold for diagonal positive-definite matrices.

V. CONCLUSION

In this work, we studied the NMF detector and, in particular, its off-grid extension. We first clarified the theoretical P_{fa} -threshold relationship, showing that it actually corresponds to the expected number of false alarms. We then analyzed the NMF performance, demonstrating that it can detect at most one target and that its performance degrades in the presence of multiple targets. Finally, we examined the impact of the noise covariance matrix and confirmed through simulations that the off-grid NMF exhibits a quasi-CFAR behavior. These findings provide a deeper understanding of the detector's capabilities and limitations in practical scenarios.

REFERENCES

- [1] S. M. Kay, *Fundamentals of statistical signal processing: estimation theory*. Prentice-Hall, Inc., 1993.
- [2] L. L. Scharf and C. Demeure, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*, ser. Addison-Wesley Series in Electrical and Computer Engineering. Addison-Wesley Publishing Company, 1991.
- [3] E. Conte, M. Lops, and G. Ricci, "Adaptive matched filter detection in spherically invariant noise," *IEEE Signal Processing Letters*, vol. 3, no. 8, pp. 248–250, 2002.
- [4] S. Kraut, L. L. Scharf, and L. T. McWhorter, "Adaptive subspace detectors," *IEEE Transactions on Signal Processing*, vol. 49, no. 1, pp. 1–16, 2001.
- [5] O. Rabaste, J. Bosse, and J.-P. Ovarlez, "Off-grid target detection with normalized matched subspace filter," in *24th European Signal Processing Conference (EUSIPCO)*, 2016, pp. 1926–1930.
- [6] P. Develter, J. Bosse, O. Rabaste, P. Forster, and J.-P. Ovarlez, "On the false alarm probability of the normalized matched filter for off-grid target detection," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 5782–5786.
- [7] O. Besson, "Adaptive detection with bounded steering vectors mismatch angle," *IEEE Transactions on Signal Processing*, vol. 55, no. 4, pp. 1560–1564, 2007.
- [8] S. Trottier, J. Bosse, O. Rabaste, P. Forster, and J.-P. Ovarlez, "Off-grid radar detection strategy for the normalized matched filter: achieving the glrt performance," in *33rd European Signal Processing Conference (EUSIPCO)*, 2025.
- [9] P. Develter, J. Bosse, O. Rabaste, P. Forster, and J.-P. Ovarlez, "On the false alarm probability of the normalized matched filter for off-grid targets: A geometrical approach and its validity conditions," *IEEE Transactions on Signal Processing*, vol. 72, pp. 982–996, 2024.
- [10] H. Weyl, "On the volume of tubes," *American Journal of Mathematics*, vol. 61, no. 2, pp. 461–472, 1939.
- [11] A. Gray, *Tubes*. Springer Science & Business Media, 2003, vol. 221.
- [12] E. Ollila, D. E. Tyler, V. Koivunen, and H. V. Poor, "Complex elliptically symmetric distributions: Survey, new results and applications," *IEEE Transactions on Signal Processing*, vol. 60, no. 11, pp. 5597–5625, 2012.