



HAL
open science

Robust control of varying weak hyperspectral target detection with sparse non-negative representation

Raphael Bacher, Céline Meillier, Florent Chatelain, Olivier J.J. Michel

► To cite this version:

Raphael Bacher, Céline Meillier, Florent Chatelain, Olivier J.J. Michel. Robust control of varying weak hyperspectral target detection with sparse non-negative representation. *IEEE Transactions on Signal Processing*, 2017, 65 (13), pp.3538-3550. 10.1109/TSP.2017.2688965 . hal-01496201

HAL Id: hal-01496201

<https://hal.science/hal-01496201v1>

Submitted on 29 Mar 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Robust control of varying weak hyperspectral target detection with sparse non-negative representation

Raphael Bacher, Celine Meillier, Florent Chatelain and Olivier Michel

Abstract—In this study, a multiple-comparison approach is developed for detecting faint hyperspectral sources. The detection method relies on a sparse and non-negative representation on a highly coherent dictionary to track a spatially varying source. A robust control of the detection errors is ensured by learning the test statistic distributions on the data. The resulting control is based on the false discovery rate, to take into account the large number of pixels to be tested. This method is applied to data recently recorded by the three-dimensional spectrograph Multi-Unit Spectrograph Explorer.

I. INTRODUCTION

WITH the constant development of new imaging devices, the exploration of massive multi-modalities datasets has become an important field of study, with many challenges to face. In particular, the present study is motivated by the need to detect faint emission line features in massive hyperspectral data produced by the recent three-dimensional (3D) spectrograph Multi-Unit Spectrograph Explorer (MUSE) instrument [1]. The targeted emission lines are markers of spatially extended structures (or ‘halos’, surrounding galaxies) that can exhibit spectral variability (spectral shifts). Furthermore, the presence of a large number of nuisance sources with high dynamics in the background makes the estimation of the background statistics particularly challenging. Searching faint signals in massive datasets requires the removal of possibly strong contributions of unwanted objects. Typically, for tracking faint signatures in hyperspectral data with high background dynamics, the spectra baseline must be estimated and removed. Finally, the size of the data to be explored calls for error controls of mis-detections with a global significance, such as the false discovery rate (FDR)[2], to allow unsupervised detection of the targets.

To summarize, the present study addresses a quite general detection problem whose main features are:

- highly variable, partially known, weak target signatures;
- background difficult to model;
- possible presence of strong contributions from unwanted sources that have to be removed;
- necessity for robust global mis-detections control.

Many methods have been developed in recent years to detect targets in hyperspectral data [3], all requiring knowledge of the background and/or the target signature. Among these studies, spectral anomaly detectors can be used when the searched signal is unknown, but rely on a parametric statistical

modeling of the background signature [4] Spectral matching detectors, such as adaptive matched filters [5] or adaptive cosine estimators [6], exploit prior knowledge of both the target signature and the background characterization. A common feature of these approaches is the lack of global control of false alarm rate. Furthermore, these detectors were developed in a remote-sensing framework [7] at rather high signal-to-noise ratio (SNR) and can barely be adapted to the new challenges (low SNR, absence of ground-truth...) raised by new massive hyperspectral data such as those provided by new instruments in astronomy. An alternate class of methods relies on sparse representation [8] based on a trained dictionary. These are in fact mostly reconstruction methods, exploited for detection purposes. Up to our knowledge these methods do not allow to calibrate the type I error (false alarm) control of the test. Another pitfall is that in general no training set are available in the astronomical context. Finally, recent approaches [9], [10] try to tackle a problem having the same features as ours. However, the Generalized Likelihood Ratio based solutions proposed in these studies do not allow reliable control of detection error. This control is crucial for massive datasets in general, and for the present MUSE dataflow in particular, and is the core of this study.

Rare events or sparse signals detection problems have received much attention in the recent statistical literature. Multiple-comparison procedures, such as higher criticism or Bonferroni-type methods, have been proposed and shown to have asymptotic optimal detection properties under sparsity regimes [11], [12], [13], [14]. Such methods do not require specification of the signals/events to be detected. Higher criticism procedures can be viewed as adaptive to the unknown sparsity level and power of the signals to be detected. These multiple-comparison procedures can be applied on the dictionary-based representations of the signals. Overcomplete and/or coherent dictionaries are prone to provide sparse representation matched to the application at hand, and were shown to greatly improve the detection power[11]. Again no global error control is performed by these methods.

The purpose of the present paper is thus threefold:

- derive a detection algorithm that benefits from the detection power of the sparse representation based approaches;
- propose a method that requires very weak assumptions on the background or target statistical properties;
- operate a test procedure that allow a control of the global false alarm rate (FDR).

The proposed method is built upon a spectral matching approach over a highly coherent dictionary of target spectra,

The authors are with the Image and Signal Department, GIPSA-Lab, Grenoble Institute of Technology, Saint Martin d’Heres 38400, France (e-mail: raphael.bacher@gipsa-lab.grenoble-inp.fr).

to take the variability into account. It elaborates on the max-test study presented in [13] for the detection of rare events. As such, it is built on a sparse representation whose purpose is to allow the formulation of a detection test, that do not need any signal reconstruction. To insure a calibration of the test procedure (FDR wise) that is robust to background misspecification, a new simple procedure is proposed, that mainly requires symmetry of the noise distribution. Note that exploiting this symmetry of the noise versus the positivity of sources in astronomical context was also developed in [15] but without providing a global control of the errors nor the formalization of a varying target matched over a dictionary of spectral shapes.

The paper is organized as follows. Section II defines the core of the proposed detection approach. The application-oriented design of the dictionary, the data preprocessing step, and the results on the real MUSE data are described in section III. Some conclusions and perspectives are drawn up in section IV.

Notations

In the following, a 'pixel' refers indifferently to a position in the MUSE spatial grid and to the associated spectrum. A pixel or its associated spectrum vector is represented by bold letters e.g. \mathbf{x} , and bold capital letters refer to matrices.

II. DETECTION METHOD

A. Testing problem

We address now the detection of a signal \mathbf{x} , from noisy data $\mathbf{y} \in \mathbb{R}^l$. Let \mathcal{H}_0 and \mathcal{H}_1 be the hypotheses denoting, respectively, the absence or presence of the source contribution \mathbf{x} . The testing problem is:

$$\begin{cases} \mathcal{H}_0 : \mathbf{y} = \boldsymbol{\epsilon}, \\ \mathcal{H}_1 : \mathbf{y} = \mathbf{x} + \boldsymbol{\epsilon}, \end{cases} \quad (1)$$

where $\boldsymbol{\epsilon} \in \mathbb{R}^l$ is a noise vector, centered and independent of \mathbf{x} , for which the distribution is not known.

When \mathbf{x} is not fully specified, a classical approach for (1) consists of modeling \mathbf{x} as a sparse superposition of reference signals taken from a massively overcomplete dictionary \mathbf{D} , see for instance [16], or [17] for classification tasks. The reference signals, or *atoms*, correspond to the column vector $\mathbf{d}_j \in \mathbb{R}^l$, for $1 \leq j \leq m$, of $\mathbf{D} \in \mathbb{R}^{l \times m}$, where m is the total number of references. These atoms are usually scaled to be ℓ_2 -normalized: $\|\mathbf{d}_j\|_2 = 1$ for $1 \leq j \leq m$.

Moreover, in the present context, the signal of interest \mathbf{x} is generally assumed to be non-negative. Thus, to enforce a non-negative decomposition, the atoms \mathbf{d}_j , for $1 \leq j \leq m$ are assumed to be non-negative. It should be noted that, unlike most of the sparse representation techniques in the literature, we do not seek to build an optimal dictionary for reconstruction/estimation but for the design of a good detection test. The dictionary construction, based on physical priors and specific to the application at hand, will be addressed in section III-C1 in the framework of MUSE application¹.

¹Note that in this application the non-negativity constraints can be relaxed. The target \mathbf{x} and the atoms \mathbf{d}_j can have negative or positive contributions, as long as their dot product $\mathbf{x}^T \mathbf{d}_j$ are non-negative for $1 \leq j \leq m$.

Under the non-negativity and sparsity assumptions, the target signal can be expressed as

$$\begin{aligned} \mathbf{x} &\approx a_{i_1} \mathbf{d}_{i_1} + \dots + a_{i_k} \mathbf{d}_{i_k}, \\ \text{s.t. } a_{i_j} &> 0, \quad \text{for } 1 \leq j \leq k, \end{aligned} \quad (2)$$

where $k \ll m$. Based on this representation, the detection problem reduces to a multiple comparison procedure with one-sided tests:

$$\begin{cases} \mathcal{H}_0 : a_1 = a_2 = \dots = a_m = 0, \\ \mathcal{H}_1 : \text{at least one } a_i > 0, \end{cases} \quad (3)$$

B. Test statistic

We now search a test statistic adapted to the detection problem (3) obtained by sparse representation. Let us first consider the statistic for a single atom. Let $S(\mathbf{y}, \mathbf{d})$ be a measure of similarity between the observed data $\mathbf{y} \in \mathbb{R}^l$ and a normalized reference vector $\mathbf{d} \in \mathbb{R}^l$. Popular examples of similarity measures include the matched filter statistic

$$S(\mathbf{y}, \mathbf{d}) \equiv \left\langle \frac{\mathbf{d}}{\|\mathbf{d}\|}, \mathbf{y} \right\rangle = \mathbf{d}^T \mathbf{y}, \quad (4)$$

or the spectral angular distance (SAD)

$$S(\mathbf{y}, \mathbf{d}) \equiv \frac{\langle \mathbf{d}, \mathbf{y} \rangle}{\|\mathbf{d}\| \|\mathbf{y}\|} = \frac{\mathbf{d}^T \mathbf{y}}{\|\mathbf{y}\|}, \quad (5)$$

which is a classical distance in hyperspectral analysis [18]. For a given signal amplitude $a = \|\mathbf{y}\| > 0$, such similarity measures are maximized when $\mathbf{y} = a\mathbf{d}$.

Based now on the pairwise similarity measures $S(\mathbf{y}, \mathbf{d}_j)$ between the observation \mathbf{y} and the dictionary atoms \mathbf{d}_j , for $1 \leq j \leq m$, a global test for the multiple (with regards to the m atoms) testing problem introduced in (3) can be derived from a Bonferroni-like correction. Accounting for (2), this leads us to consider the following one-sided max-test approach:

$$T_{\max}(\mathbf{y}) \equiv \max_{1 \leq j \leq m} S(\mathbf{y}, \mathbf{d}_j) \underset{H_0}{\overset{H_1}{\geq}} \eta, \quad (6)$$

where η is a given threshold. The motivations for using this max-test approach are two-fold. First, from a theoretical point of view, with highly sparse signals, the max-test method is asymptotically as efficient as the asymptotically optimal higher criticism method, as demonstrated in [11], [13]. Secondly, in finite sample settings such as the MUSE hyperspectral datasets, max-test has been shown to be relatively efficient [9], [19], and empirically more powerful than higher criticism methods [20].

We now tackle the problem of applying the max-test defined in (6) to a large number n of data realisations $\{\mathbf{y}_i\}_{1 \leq i \leq n}$. We are again in a multiple-testing context, now with regards to the number of samples n . This is in regards to this context that we seek to control the detection errors. To fix the decision threshold η while controlling the type I errors, i.e., the samples under \mathcal{H}_0 that will be falsely detected as \mathcal{H}_1 , the distribution under the null hypothesis \mathcal{H}_0 of the max-test statistic $T_{\max}(\mathbf{y})$ must be known. In real applications such as the MUSE data, due to the physical process and preprocessing steps (e.g.,

interpolation, background subtraction), noise is spatially and spectrally correlated with an unknown complex dependence structure. Thus the distribution of $T_{\max}(\epsilon)$, where ϵ is the noise vector introduced in (1), cannot be easily modeled as a standard parametric distribution. However in a large scale testing framework, it becomes possible to estimate this distribution, as explained in the next section.

C. Learning the null distribution

Consider now the following assumptions:

A1. *The noise vector ϵ is centered and symmetrically distributed: ϵ and $-\epsilon$ have the same distribution,*

A2. *The similarity measure $S(\mathbf{y}, \mathbf{d})$ used to construct the max-test statistics is an odd function of the observations \mathbf{y} , i.e. $S(\mathbf{y}, \mathbf{d}) = -S(-\mathbf{y}, \mathbf{d})$ for any sample \mathbf{y} and for any reference vector \mathbf{d} .*

Assumption A1 on the noise is relatively weak and fairly reasonable. This is, for instance, satisfied for any centered elliptical distribution, such as multivariate Gaussian or student distributions. Note also that assumption A2 is clearly satisfied for the matched filter or the SAD statistics described respectively in (4) and (5). A direct consequence of these assumptions is the following key property:

Proposition II.1. *Based on assumptions A1 and A2, the max-test statistic $T_{\max}(\mathbf{y})$ and the opposite of the min statistic $-T_{\min}(\mathbf{y})$, where $T_{\min}(\mathbf{y}) \equiv \min_j S(\mathbf{y}, \mathbf{d}_j)$, are identically distributed under the null hypothesis \mathcal{H}_0 .*

Proof. Under the null hypothesis $\mathbf{y} = \epsilon$. According to A1, $T_{\max}(\epsilon)$ and $T_{\max}(-\epsilon)$ are identically distributed. Moreover,

$$\begin{aligned} T_{\max}(-\epsilon) &= \max_j S(-\epsilon, \mathbf{d}_j) = -\min_j \{-S(-\epsilon, \mathbf{d}_j)\}, \\ &= -\min_j S(\epsilon, \mathbf{d}_j) = -T_{\min}(\epsilon), \end{aligned}$$

where the first equality on the second line is due to A2. Thus $T_{\max}(\epsilon)$ and $-T_{\min}(\epsilon)$ have the same distribution. \square

In a large-scale testing framework (with regards to the number of samples n), the max and min statistics $T_{\max}(\mathbf{y})$ and $T_{\min}(\mathbf{y})$ are computed for a large number of observations \mathbf{y}_i , for $1 \leq i \leq n$. Let $\pi_0 \in (0, 1]$ be the true proportion of observations \mathbf{y}_i distributed according to the null hypothesis, while $\pi_1 = 1 - \pi_0$ is the proportion of observations distributed according to the alternative hypothesis for the testing problem (1). Let $F(t) = \Pr(T_{\max}(\mathbf{y}) \leq t)$ be the cumulative distribution function of the max statistic $T_{\max}(\mathbf{y})$. This distribution function can be expressed as a two-groups model:

$$F(t) = \pi_0 F_0(t) + \pi_1 F_1(t),$$

where F_0 and F_1 denote the distribution functions under the null and the alternative hypotheses, respectively. Under the non-negativity assumption introduced in section II-B, $T_{\max}(\mathbf{y})$ should be stochastically larger under \mathcal{H}_1 than under \mathcal{H}_0 , i.e., $F_0(t) > F_1(t)$ for any $t \in \mathbb{R}$. Let μ_0 be the median of the

max test statistics under the null hypothesis² (referred to as the *null median* hereafter), i.e., $F_0(\mu_0) = \frac{1}{2}$. We now introduce the classical *zero assumption*, as termed by Efron a different context [21, Chap. 6]. This assumes the existence of a noise-only domain that allows to build a procedure for estimating the null distribution (see *Remark 2* hereinafter for a discussion about this assumption).

A3 (Zero assumption for F_1). $F_1(t) = 0$ for the region $t \leq \mu_0$ where the max statistics are most likely under \mathcal{H}_0 .

From this assumption, we can now derive the following expression:

$$F(t) = \pi_0 F_0(t), \quad \text{for } t \leq \mu_0.$$

In a similar manner, the survival function $G(t) = \Pr(-T_{\min}(\mathbf{y}) > t)$ of the opposite min statistic $-T_{\min}(\mathbf{y})$ reads as

$$G(t) = \pi_0 G_0(t) + \pi_1 G_1(t),$$

where G_0 and G_1 are the survival functions of $-T_{\min}(\mathbf{y})$ under the null and alternative hypotheses, respectively. This comes from the non-negativity assumption that $-T_{\min}(\mathbf{y})$ should be stochastically smaller under \mathcal{H}_1 than under \mathcal{H}_0 , i.e., $G_0(t) > G_1(t)$. Note that μ_0 is also the median of the null distribution of $-T_{\min}(\mathbf{y})$ as $G_0(\mu_0) = 1 - F_0(\mu_0) = \frac{1}{2}$ according to the proposition II.1. This allows us to introduce the following zero assumption.

A4 (Zero assumption for G_1). $G_1(t) = 0$ for the region $t \geq \mu_0$ where the opposite min statistics are most likely under \mathcal{H}_0 .

Thus

$$G(t) = \pi_0 G_0(t), \quad \text{for } t \geq \mu_0.$$

Since $G_0(t) = 1 - F_0(t)$ according to the proposition II.1, we can finally derive the following expression:

$$\pi_0 F_0(t) = \begin{cases} F(t), & \text{for } t \leq \mu_0, \\ \pi_0 - G(t), & \text{for } t > \mu_0. \end{cases} \quad (7)$$

The main interest of this expression is that it does not depend on each group distribution function but only on the distribution function of the two-groups model. In particular, assumptions A3 and A4 do not require to fully specify F_1 , which is unlikely to be known in practice. Expression (7) is therefore robust to alternative miss-specifications. This expression still depends on the theoretical null median μ_0 , the proportion π_0 of samples under \mathcal{H}_0 , and the distribution functions $F(t)$ and $G(t)$, which are not known. However, when a large number of observations $\mathbf{y}_1, \dots, \mathbf{y}_n$ are available, these quantities can be estimated from the empirical distributions of the test statistics. Let

$$\bar{F}(t) = \frac{\#\{T_{\max}(\mathbf{y}_i) \leq t\}}{n}, \quad \bar{G}(t) = \frac{\#\{-T_{\min}(\mathbf{y}_i) > t\}}{n},$$

be the empirical distribution function of T_{\max} and the empirical survival function of $-T_{\min}$, respectively.

²For the sake of simplicity, the observations are assumed to obey absolutely continuous distributions. Thus the test statistics T are also continuous, and their median μ is defined as $\Pr(T \leq \mu) = \Pr(T \geq \mu) = \frac{1}{2}$. The extension to discrete statistics is left to the reader.

A5 (Weak dependence assumption). *The empirical functions $\overline{F}(t)$ and $\overline{G}(t)$ converge uniformly toward the theoretical distribution functions $F(t)$ and $G(t)$, respectively:*

$$\sup_t |\overline{F}(t) - F(t)| \rightarrow 0, \quad \sup_t |\overline{G}(t) - G(t)| \rightarrow 0,$$

almost surely as the number n of observations grows to infinity.

Statement A5 is verified under weak dependence conditions between the observed samples $\mathbf{y}_1, \dots, \mathbf{y}_n$. In particular, this holds for independent or short-range dependent samples according to the Glivenko-Cantelli theorem. A direct consequence is the pointwise convergence in probability of $\overline{F}(t)$ and $\overline{G}(t)$ toward $F(t)$ and $G(t)$, respectively, for any $t \in \mathbb{R}$.

Due to zero assumptions A3 and A4 and proposition II.1, μ_0 satisfies $F(\mu_0) = G(\mu_0) = \frac{\pi_0}{2}$. Therefore, based on assumption A5, an estimator of the null median μ_0 can be searched as a solution of the following equation for μ :

$$\overline{F}(\mu) = \overline{G}(\mu). \quad (8)$$

Lemma II.2 (Empirical null median estimator). *Let $t_{(1)} < t_{(2)} < \dots < t_{(2n)}$ be the ordered values of the statistics belonging to the sample $\mathbf{t} = (T_{\max}(\mathbf{y}_1), \dots, T_{\max}(\mathbf{y}_n), -T_{\min}(\mathbf{y}_1), \dots, -T_{\min}(\mathbf{y}_n))$. Let $\hat{\mu}_0$ be the sample median of \mathbf{t} , which is defined as*

$$\hat{\mu}_0 = \frac{t_{(n)} + t_{(n+1)}}{2}. \quad (9)$$

Then $\hat{\mu}_0$ satisfies (8) and is a consistent estimator of the null median μ_0 , under zero assumptions A3 and A4.

Proof. See appendix A1. \square

Based on the null median estimator given in lemma II.2, we can now obtain the empirical null estimates of π_0 and $F_0(t)$. Let

$$\mathbf{s}_0 = \{T_{\max}(\mathbf{y}_i) \leq \hat{\mu}_0\},$$

be the sample set of the max-test statistics truncated on $(-\infty, \hat{\mu}_0]$, the elements of which are denoted as $s_{0,i}$, for $1 \leq i \leq n_0$, and where $n_0 = |\mathbf{s}_0|$. Similarly,

$$\mathbf{g}_0 = \{-T_{\min}(\mathbf{y}_i) > \hat{\mu}_0\},$$

denotes the set of the opposite min statistics truncated on $(\hat{\mu}_0, +\infty)$, the elements of which are denoted as $g_{0,i}$ for $1 \leq i \leq n_0$ (according to lemma II.2, these two sets are of equal size).

Proposition II.3 (Empirical estimators under \mathcal{H}_0). *Under assumptions A3 and A4,*

$$\hat{\pi}_0 = \min \{2n_0/n, 1\}, \quad (10)$$

is a consistent estimator of the null proportion π_0 , and

$$\hat{F}_0(t) = \frac{\#\{s_{0,i} \leq t\} + \#\{g_{0,i} \leq t\}}{2n_0} \quad (11)$$

is a pointwise consistent estimator of the null distribution $F_0(t)$, for $t \in \mathbb{R}$.

Proof. See appendix A2. \square

Remark 1: the dependence structure across a set of observations $\mathbf{y}_1, \dots, \mathbf{y}_n$ with $\mathbf{y}_i \in \mathbb{R}^l$ is not required to specify the empirical estimators $\hat{\pi}_0$ and \hat{F}_0 given in proposition II.3. These non-parametric estimates rely essentially on the noise symmetry assumption A1, which is very weak. As a consequence, these estimators are robust to miss-specifications that are prone to occur with parametric assumptions.

Remark 2: Zero assumptions A3 and A4 provide an idealized mathematical framework for which the empirical estimators are shown to be consistent. However, these assumptions are unlikely to be satisfied in practice. As a consequence, Eq. (7) is an approximation. Note, however, that the closer π_0 is to one, the more accurate the approximation is. This is the case in many large-scale testing problems, where the null proportion π_0 is usually close to one. The approximation gains also in accuracy the more $F_0(t)$ (resp. $G_0(t)$) dominates $F_1(t)$ (resp. $G_1(t)$) for $t \leq \mu_0$ (resp. for $t \geq \mu_0$). Moreover, if a few observations distributed according to the alternative distribution belong to the regions where they are assumed to be absent, then $\hat{\pi}_0$ tends to be biased upward, and $\hat{F}_0(t)$ tends to be slightly biased toward the alternative distribution $F_1(t)$. It is of note that from a statistical testing perspective, this slight bias goes in the good way. Indeed, a detection procedure based on $\hat{F}_0(t)$ (and possibly $\hat{\pi}_0$) then becomes more conservative as p -values are slightly biased upward. This results in a small loss of power but the control of type I errors is still (asymptotically) guaranteed.

Figure 1 shows the empirical density functions associated with the max-test statistics $T_{\max}(\mathbf{y}_i)$ and the opposite min statistics $-T_{\min}(\mathbf{y}_i)$ for synthetic data with a testing framework that mimics the MUSE one. We can see that the max density has a heavier right tail than the opposite min one. This is due to the contribution of the \mathcal{H}_1 samples, while the opposite min density right tail is (approximately) distributed according to the theoretical null density due to (approximation) A4. By symmetry, the opposite min density has a heavier left tail than the max density one.

To appreciate the accuracy of the empirical estimators given in proposition II.3, the empirical null-density function that is obtained from the right-truncated sample \mathbf{s}_0 and the left-truncated sample \mathbf{g}_0 that are used to construct $\hat{F}_0(t)$, is depicted in Figure 2a. Here, the null proportion estimate is larger than the theoretical value: $\hat{\pi}_0 = 0.89$ and $\pi_0 = 0.81$. Nevertheless, the empirical null-density function is very close to the theoretical one. This is confirmed by the quantile-quantile plot between \hat{F}_0 and the theoretical distribution F_0 shown in Figure 2b. In particular, this remains true in the distributions tails, where the accuracy of the quantile estimates is crucial to the robustness of the test at low control levels.

D. Error control

In multiple testing (around $n = 2500$ tested pixels for a 50×50 patch in the MUSE context), the classical Type I error control of each individual test might not be appropriate; see e.g., [22], [23]. Indeed the number of wrongly rejected null hypotheses can become relatively important (i.e., even larger than the number of true detections) due to the high




Fig. 1: Empirical density functions of the max-test statistics $T_{\max}(\mathbf{y}_i)$ (blue line with \blacklozenge markers) and of the opposites of min statistics $-T_{\min}(\mathbf{y}_i)$ (green line) for $n = 2500$ independent samples $\mathbf{y}_1, \dots, \mathbf{y}_n \in \mathbb{R}^l$ generated from the observation model (1) with $l = 30$. Noise vector ϵ_i has i.i.d. entries that follow a student distribution with $\nu = 5$ degrees of freedom. The proportion of null hypotheses is $\pi_0 = 0.81$. Samples under \mathcal{H}_1 are generated as $\mathbf{y}_i = a_i \mathbf{d} + \epsilon_i$, where $a_i \in [0.1, 3]$ and \mathbf{d} has unit length. The min and max-test statistics are obtained from a dictionary \mathbf{D} with $m = 15$ positively correlated atoms, including \mathbf{d} , and using the SAD similarity measure.

number of tests. To address this kind of issue, a global error control approach, namely the FDR, was introduced in [2]. The FDR controls the expected proportion of true null hypotheses wrongly rejected, which are referred to as the *false discoveries*, among all of the rejected tests:

$$\text{FDR} = E \left[\frac{U}{\max(R, 1)} \right],$$


where R is the total number of tests where the null hypothesis is rejected, while U is the number of false discoveries among the R discoveries. A simple and widely used approach to control this FDR is the Benjamini and Hochberg (BH) procedure that was also developed in [2]. Let p_i be the p -value associated with the i th test statistics. Let $p_{(1)} \leq p_{(2)} \leq \dots \leq p_{(n)}$ now be the ordered p -values, and $\mathcal{H}_0^{(1)}, \dots, \mathcal{H}_0^{(n)}$ denote the null hypotheses for this ordering. For a preselected control level $0 \leq q \leq 1$, the BH_q procedure rejects $\mathcal{H}_0^{(1)}, \dots, \mathcal{H}_0^{(\hat{k})}$ where

$$\hat{k} = \max \left\{ 0 \leq k \leq n : p_{(k)} \leq q \frac{k}{n} \right\},$$


with $p_{(0)} = 0$ by convention. In our right-sided detection framework, the i th empirical p -value, can be derived from the empirical null distribution given in proposition II.3 as

$$p_i = 1 - \hat{F}_0(T_{\max}(\mathbf{y}_i)), \quad \text{for } 1 \leq i \leq n. \quad (12)$$

Then in the case of n independent tests, or under specific positive dependencies [24], the BH_q procedure controls the FDR at a level $\pi_0 q \leq q$. Thus, if π_0 is known, the BH



(a) Empirical density function of the sample set $\mathbf{t}_0 = \mathbf{s}_0 \cup \mathbf{g}_0$, where \mathbf{s}_0 and \mathbf{g}_0 are the sample sets used to construct $\hat{F}_0(t)$ (blue curve) and theoretical density function of $T_{\max}(\mathbf{y})$ under \mathcal{H}_0 (green curve with \blacklozenge marker).



(b) QQ-plot of $\hat{F}_0(t)$ against the theoretical distribution F_0 , \bullet marker, and $y = x$ line (red dashed line).

Fig. 2: Comparisons between the empirical null distribution estimator $\hat{F}_0(t)$ and the theoretical distribution of $T_{\max}(\mathbf{y})$ under \mathcal{H}_0 . Data are generated using the same setting as in Figure 1. The theoretical distribution function was estimated using 10^5 Monte-Carlo runs.

procedure can be applied at the nominal level $\frac{q}{\pi_0}$ to improve its power while controlling FDR at level q . Building on this idea, Storey [25] proposed the following *modified* estimator of the null proportion π_0 :

$$\hat{\pi}_0^*(\zeta) = \min \left\{ \frac{1 + \#\{p_i > \zeta\}}{(1 - \zeta)n}, 1 \right\}, \quad \text{for } \zeta \in [0, 1),$$

where ζ has to be arbitrarily fixed (usually at $\frac{1}{2}$). Storey showed that under the weak dependence assumption A5, the $\text{BH}_{q'}$ procedure at nominal level $q' = q/\hat{\pi}_0^*(\zeta)$ asymptotically controls the FDR at level q . We show in the following that the

same strategy can be applied with the empirical null-estimator $\hat{\pi}_0$.

Proposition II.4 (Storey π_0 estimator). *The empirical null estimator $\hat{\pi}_0$ defined in (10) and the Storey estimator $\hat{\pi}_0^*(\zeta)$ derived from the empirical p -values defined in (12) are equal for any $\zeta = \frac{k}{2n_0}$ with $k \in \{n_0, \dots, 2n_0 - 1\}$, and are asymptotically equivalent for any $\zeta \in [\frac{1}{2}, 1)$.*

Proof. See appendix B. \square

This equivalence is not surprising since, like the proposed empirical null estimators, Storey estimator is based on a zero assumption (i.e., the p -value density function under \mathcal{H}_1 is zero on $(\zeta, 1]$).

This leads us to consider the following multiple testing procedure described in Alg. 1.

Algorithm 1 FDR-based detection procedure

Input: nominal FDR level q

- 1) compute the empirical null estimators $\hat{\pi}_0$ and \hat{F}_0 as defined in proposition II.3;
 - 2) compute the empirical p -values according to (12);
 - 3) apply the BH procedure at a nominal control level $q/\hat{\pi}_0$.
-

Note that it has been shown in [22] that for matched filter statistics, with a non-negative template and under Gaussian assumptions, the test statistics obey a positive regression dependence on a subset (PRDS) condition. Therefore, the BH procedure ensures exact FDR control in finite sample settings. Here the problem is more complex. The test statistics are derived from extreme values that can be correlated. Then PRDS is difficult to ensure theoretically, even under Gaussian assumptions on the noise. However, under weak dependence assumption A5, an Oracle procedure similar to Alg. 1, but where the p -values are computed from the theoretical null distribution F_0 , can be proven to control asymptotically the FDR at level q [25]. As discussed in the previous subsection, the p -value empirical estimates tend to be slightly biased in a conservative way. Moreover, if the null distribution can be estimated on a larger sample than the sample to be tested, the variance of these estimates can be reduced. This supports the asymptotic control of the proposed procedure.

E. Validation

Figure 3 shows the FDR obtained with Alg. 1, on 3D (spatial + spectral) simulated data that are subjected to weak dependence (spatial kernel convolution), for different levels of nominal control q and different signal-to-noise ratio (SNR). The SNR is defined here as $10 \log \frac{A}{nl\sigma^2}$, where n is the number of pixels (i.e., the number of tests to perform), l is the number of spectral bands (i.e., the dimension of the observations \mathbf{y}_i), σ^2 is the marginal variance of the noise, and $A = \|\mathbf{x}\|^2$ is the energy of the 3D contribution of the signals to be detected. The experimental set-up is similar to Figure 1, except that a spatial convolution kernel of size 3 by 3 was applied to create local spatial correlations. The empirical null estimators defined on section II-C were computed from extended cubes of 200 by 200 pixels by 30 wavelengths.

This figure emphasizes that control of the FDR is correctly achieved for the different SNR levels. As expected, due to the zero assumption approximations, Alg. 1 is a little more conservative than the Oracle one (based on the true F_0 and π_0) at low SNR, where the alternative distribution is closer to the null one.




Fig. 3: Empirical FDR (averaged from 1000 Monte-Carlo runs) versus the SNR under weak dependence, for different levels of FDR control $q = 0.02, 0.05, 0.1, 0.2$ (in cyan, red, green, and blue, respectively). Dashed horizontal lines, nominal control levels q ; \blacklozenge curves, FDR for Alg. 1 empirical procedure; \star curves, FDR for Oracle procedure based on the true (computed from 10^5 Monte-Carlo runs) F_0 and π_0 . Data are generated as 3D cubes of $n = 51 \times 51$ pixels by $l = 30$ wavelengths.

Table I shows the advantage of assuring a global control with a detection threshold that adapts to the data. It compares this global control with a pixel-wise control based on the probability of false alarm (PFA). Controlling with a η (e.g. 5%) PFA threshold results in detecting all pixels with p -values smaller than η . Such a PFA control then ensures that in average a fraction η of all the tested pixels will be wrongly detected (but says nothing of the proportion of these wrong detections among the detected set).

When confronted to noise-only data, a control procedure with PFA at level 5% detects 144 spurious pixels, that is the size of a possible source. To insure that no source is falsely detected, we may turn to a more conservative level, e.g. 0.1%; this results in poor source detection power (around 55%) whereas a 5% level led to very good power (82%) at the price of a large number of false alarms (false discovery proportion $\simeq 43\%$). On the same dataset, a FDR control at 20% does not lead here to any false detection in the absence of source while maintaining a high detection power (72%) in presence of a source, thus adapting to the data. Table I shows that the false discovery proportion is around 18% for a nominal FDR control of 20%. Note that by applying our procedure we can estimate the FDR level that matches a given detection threshold. For instance a threshold on the p -values associated with a 5% PFA level yields a FDR estimate around 44% on data tested here.

TABLE I: Comparison between FDR and PFA control on data with and without target. Data was built from noise-only regions in MUSE real data and a synthetic source was added. Number of tests is 2500 (50 by 50 pixels) and source size is 185 pixels. Results were averaged on 5 different regions.

	PFA 0.05	PFA 0.001	FDR 0.2
<i>Noise-only</i>			
False Detections (pixels)	144	2	0
True Detections (pixels)	0	0	0
<i>Source + noise area</i>			
False Detections (pixels)	117	2	30
True Detections (pixels)	153	106	133
False discovery proportion (%)	43.3	1.8	18.4

Table I shows that the false discovery proportion is indeed around 44% for a threshold corresponding to a 5% PFA.

In the next paragraph the ability of the proposed method to control error rate is compared with a generalized likelihood ratio (GLR) approach (inspired by [9] and [10]). Noise is supposed centered Gaussian $\epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma)$ where the covariance matrix $\Sigma \in \mathbb{R}^{l \times l}$ is assumed to be diagonal. We have the following detection test :

$$\begin{cases} \mathcal{H}_0 : \mathbf{y} = \epsilon, \\ \mathcal{H}_1 : \mathbf{y} = \mathbf{D}\mathbf{a} + \epsilon, \text{ with } \|\mathbf{a}\|_0 = 1, \mathbf{a} \geq \mathbf{0} \end{cases}$$

where $\|\cdot\|_0$ is the ℓ_0 -pseudo-norm (number of non-zero components) and $\mathbf{a} \geq \mathbf{0}$ is the non-negativity constraint on the coefficients. The GLR test with 1-sparsity constraint yields the following test statistic([9])

$$T_{GLR}(\mathbf{y}) = \frac{\max_{\mathbf{a}} p(\mathbf{y}|\mathbf{D}\mathbf{a}, \mathcal{H}_1)}{p(\mathbf{y}|\mathcal{H}_0)} \text{ s.t. } \|\mathbf{a}\|_0 = 1, \mathbf{a} \geq \mathbf{0},$$

where $p(\mathbf{y}|\mathbf{D}\mathbf{a}, \mathcal{H}_1)$ denotes the probability density function of \mathbf{y} under \mathcal{H}_1 and $p(\mathbf{y}|\mathcal{H}_0)$ denotes the probability density function of \mathbf{y} under \mathcal{H}_0 . Using the Gaussian assumption it comes that

$$T_{GLR}(\mathbf{y}) = \frac{\mathbf{d}_{\hat{j}}^T \hat{\Sigma}^{-1} \mathbf{y}}{\sqrt{\mathbf{d}_{\hat{j}}^T \hat{\Sigma}^{-1} \mathbf{d}_{\hat{j}}}}$$

where \hat{j} is the index of the non-zero component of the optimal $\hat{\mathbf{a}}$ for the GLR statistics, and $\hat{\Sigma}$ is estimated from the residuals. There is no closed-form expression of the distribution of this statistic since it consists in taking the max of a correlated Gaussian vector. Thus we calibrated this statistic under \mathcal{H}_0 (normal centered noise) by Monte-Carlo.

Figures 4 and 5 illustrate the main advantage of the proposed method : the control is ensured when the noise distribution is symmetrical without further assumptions. The GLR approach has to be calibrated under \mathcal{H}_0 distribution so any deviation from the theoretical \mathcal{H}_0 results in a loss of control, as illustrated by figure 5 where the noise is drawn from a Student distribution. It can be seen that BH procedure based on the theoretical \mathcal{H}_0 GLR statistics do not correctly control FDR. In a first time the effective FDR strongly exceeds the given control level (allowing GLR to be “more powerful” at a given nominal control level). In a second time it becomes too conservative. This comportment can be explained by the Gaussian fit of the Student distribution of noise: tails are

underestimated (hence the excess in FDR) whereas the bulk is overestimated (hence the loss in power in the second part). It should be noted that a classical ROC curve of the two methods would show very similar performance (same power for an effective error budget) between the two approaches but would hide the inaccuracy of error control of the GLR approach. Moreover figure 4 shows that when GLR is at its best (adequate model), the proposed method does stays really close in term of power despite its versatility.




Fig. 4: True FDR (left) and power (right) versus nominal FDR control level, between GLR and proposed method, on synthesized data with Gaussian noise. The GLR was calibrated by 10^4 Monte-Carlo runs on normal noise. Σ was estimated on the data. Results were averaged on 200 runs of simulated data cubes with $\pi_0 = 0.97$. GLR is in blue, proposed method in green, $y = x$ in dashed red. Matched filter similarity measure (4) was used for the proposed method.




Fig. 5: True FDR (left) and power (right) versus nominal FDR control level, between GLR and proposed method, on synthesized data with Student noise (4 degrees of freedom). The GLR was calibrated by 10^4 Monte-Carlo runs on normal noise. Σ was estimated on the data. Results were averaged on 200 runs of simulated data cubes with $\pi_0 = 0.97$. GLR is in blue, proposed method in green, $y = x$ in dashed red. Matched filter similarity measure (4) was used for the proposed method.

III. APPLICATION TO THE MUSE DATA

It is believed that young galaxies are often surrounded by halos of hydrogen gas, known as the circum galactic medium. The emissions from these halos can be several orders of magnitude fainter than those from the galaxies. Furthermore, the emission spectrum of the halos are composed of narrow lines, notably with the Lyman- α line. The MUSE [1] was developed to detect such emissions. It is a 3D spectrograph that can image and analyze a field of 1 arcmin² by producing a hyperspectral data cube of 300 by 300 pixels by 3600 wavelengths. Its spectral range covers the visible and near-infrared domain, from 450 nm to 930 nm.

Since the MUSE first light in January 2014, several studies [9], [26] have already been conducted on the MUSE data. The aim has been to detect faint young galaxies, which are also characterized by the presence of a powerful Lyman- α emission line in their emission spectrum. Both methods mostly assume spatially and spectrally punctual sources. As a consequence, they efficiently find the core of galaxies, but are not (yet) adapted to detect the faint extended halos.

The purpose here is to explore the vicinity of these already detected galaxies, and track the Lyman- α emission line as far away from the galaxy as possible.

A. The MUSE data

The proposed detection method is applied to the MUSE observations of the sky region known as Hubble Deep Field South (HDFS), as it was previously observed by the Hubble space telescope. The MUSE produces huge amounts of data that have to be processed by a data reduction system before they can be used for scientific analysis. In particular, a resampling process creates local correlations in all directions (spatial and spectral) between voxels that cannot be easily modelled due to the data dimensions. The data reduction system applied to the data used here is detailed in [27]. The output is a 300 by 300 pixels by 3600 wavelengths data cube that is associated with a variance cube of the same dimensions. This latter is estimated by propagating the error estimated at the captor level at each stage of the processing.

A catalog of astronomical objects in HDFS was built in [27]. About 90 of these objects are remote galaxies known as Lyman α emitters that are likely to have a halo. For each of these sources, a spatial-spectral neighborhood is defined, which is centered spatially on the galaxy center and spectrally on the emission line peak, and a subcube of 50 by 50 pixels by 30 wavelengths is extracted. This cube extraction is performed for the two following reasons. Spectrally: the signal of interest (hydrogen Lyman α emission line) is concentrated in a few wavelengths around the emission peak. Outside of this domain, galaxy spectra (used to built reference spectrum) can contain other features that are not present in the targeted hydrogen surrounding halo. Spatially: as the targeted halo is expected to stay close to the galaxy, exploring empty (only noise) remote regions would only result in a loss in power (in our global control procedure).

B. Pre-processing workflow

To deal with the MUSE data, several pre-processing steps are needed. First, the spectral continuum is robustly estimated and removed in each pixel using [28]. Then coarse reduction of the data is carried out using the variance cube provided with the data, followed by robust centering and finer reduction, slice by slice [20]. After the reduction by the variance cube, we can make the following assumption of stationarity of the noise.

A6 (Stationary noise). *The noise is stationary on each wavelength-slice of the cube.*

Spatial matched filter preprocessing:

For ground-based astronomical instruments such as MUSE, the spatial system impulse response (FSF) is mainly due to atmospheric turbulence. This FSF is measured for each observation (see [29]), and is independent of the instrumental noise. As such, we can make the following assumption:

A7. *The noise in the observation model (1) is not filtered by the FSF.*

Based on A7, the following strategy was chosen to improve the SNR: a spatial convolution with the FSF (which is modeled as a symmetrical function) is applied to each image of the wavelength axis of the data cube. It is not strictly speaking a spatial matched filter to the searched halo. Indeed, the halo extension and its intensity profile are not known. However, this greatly improves the SNR. The price to be paid is that the theoretical spatial extension of the halo is enlarged by this operation. In practice, the halo has a larger extension than the FSF, with an intensity profile that, as does the FSF, decreases quickly toward zero on its support. Therefore, this effect can be neglected in the detection results.

C. Detection

In the application on hand, we have the following assumptions:

- 1) The galaxy spatial center is already known, as well as the spectral position of the emission line in the galaxy spectrum (with e.g., the method developed in [26]);
- 2) The emission line in the halo spectra has a shape similar to the emission line in the galaxy spectrum, the continuum of which has been subtracted, but can present a shift along the spectral wavelengths;
- 3) Samples (pixels) are weakly dependent.

The second hypothesis is only partially true: in reality the redshift is not a simple spectral shift, but the composition of a shift and a dilatation. However, at the spectral resolution, the deformation can be neglected. The third assumption is justified because dependences between pixels are mainly due to interpolations during the resampling step, and these dependences have short-range effects. Moreover, in the pre-processing workflow, only the convolution by the FSF creates significant spatial dependences, once again with a short-range effect due to the finite support of the FSF. Thus, the weak dependence assumption A5 is fulfilled. Assumptions required to apply the proposed method are also fulfilled. MUSE data result from

the summing of a high number of exposures thus the noise tends to be Gaussian (and as such symmetrical) by application of the central limit theorem. The following numerous pre-processing steps (background subtraction, centering, variance reduction,...) all keep the symmetrical property of the centered noise thus ensuring that assumption A1 hold. Assumption A2 (odd similarity measure) is fulfilled by construction as we choose the SAD measure. As pointed out in *Remark 2* of II-C, assumptions A3 and A4 can't be strictly guaranteed. However, outside of this ideal framework, the key point is that equation (7) is a good enough approximation. Indeed, the targeted signal is assumed to be distinct enough from the background noise and well approximated by the dictionary (see section III-C1); thus $T_{\max}(\mathbf{y})$ will be significantly stochastically larger under \mathcal{H}_1 than under \mathcal{H}_0 for detectable signals. Moreover the galaxy and halo pixels are supposed to be in strong minority in the spatially explored region, that is π_0 is close to one, which enforces these approximations. Finally, as stressed in *Remark 2*, the approximation errors in equation (7) can only lead to a small loss in power and the control is still guaranteed (the bias is conservative as shown in figure 3 for low SNR).

For a given galaxy, a spatial-spectral neighborhood is defined, centered spatially on the galaxy center and spectrally on the emission line peak. Based on these hypotheses, the approach developed here consists of applying the following steps.

- Estimate a reference emission line spectrum by averaging the spectra of a few pixels at the center of the galaxy.
- Create a (highly) coherent family of shifted versions of this reference target signature to build a dictionary.
- Test each pixel of the defined neighborhood using the method developed in section II.

1) *Dictionary*: One of the main assumptions here is that the target signature variability can mostly be modeled as a spectral shift. Thus, the dictionary is built here by creating shifted variants of one target signature, \mathbf{d}_* . Assuming that \mathbf{d}_* comes from sampling of a continuous model $f(\cdot)$, we can define as \mathbf{d}_*^δ , the shifted vector that is obtained by sampling $f(\cdot - \delta)$. The *linearly spaced shifts* (LSS) dictionary model on an interval $[-\tau, \tau]$ is then defined for a given size m , as the dictionary \mathbf{D}^m composed of the atoms $\mathbf{d}_k = \mathbf{d}_*^{\tau_k}$, where $\tau_k = -\tau + \frac{2\tau}{m-1}k$, for $k = 0, \dots, m-1$.

The key question is then the choice of the number m of shifted versions, or in other words, the redundancy of the dictionary. To allow a study of this parameter, we place ourselves in a simplified context:

- the noise is supposed to be i.i.d. $\mathcal{N}(0, 1)$;
- the similarity measure is a spectral matched filter between a dictionary atom and the tested spectrum, as in (4);
- the reference spectrum \mathbf{d}_* is a non-negative vector with unit length, where its autocorrelation function $\Gamma(u) = \langle \mathbf{d}_*, \mathbf{d}_*^u \rangle$ is non-increasing in $|u|$, and has compact support such that $\|\mathbf{d}_*^u\| = \|\mathbf{d}_*\| = 1$, for $u \in [-\tau, \tau]$;
- the target signature \mathbf{x} is built from a translation \mathbf{d}_0^u of the reference spectrum $\mathbf{x} = a\mathbf{d}_0^u$, where $a > 0$, and u is a random shift that is uniformly distributed on $[-\tau, \tau]$.

A measure of the redundancy of a given normalized dictio-

nary \mathbf{D} can be given by its coherence, which is defined as $\mu = \max_{i \neq j} |\langle \mathbf{d}_i, \mathbf{d}_j \rangle|$. For a LSS dictionary \mathbf{D}^m , and under the aforementioned assumptions, this coherence reduces to the correlation between two consecutive atoms: $\mu = \langle \mathbf{d}_j, \mathbf{d}_{j+1} \rangle$, for $1 \leq j < m$. As illustrated by figure 6, by design of the dictionary, the larger the dictionary size m , the more correlated the atoms are, and the more coherent the dictionary is.




Fig. 6: Example of dictionaries built from a reference \mathbf{d}_* , with varying number of atoms (with 3 atoms and 0.2 coherence and with 5 atoms and 0.5 coherence). The reference $\mathbf{d}_* \in \mathbb{R}^l$, with $l = 30$, is sampled for $j = 1, \dots, l$ from a Gaussian density centered on the median band $j = 15$, with full width at half maximum of 5 ($\sigma \approx 2.12$) truncated on ± 6 around the mode, and ℓ_2 -normalized. The maximal shift is $\tau = 8$.

Let $\mathbf{z}^m = (\mathbf{D}^m)^T \mathbf{y} \in \mathbb{R}^m$ be the vector of the matched filter statistics, the elements of which are defined as z_j^m for $1 \leq j \leq m$. For a given decision threshold η , the PFA for the max-test approach is expressed as

$$\alpha_m = \Pr(\max \mathbf{z}^m > \eta), \quad \text{under } \mathcal{H}_0. \quad (13)$$

Here the noise vector ϵ is $\mathcal{N}(0, \mathbf{I}_m)$ distributed under \mathcal{H}_0 . If the atoms are orthogonal (e.g., if they have disjoint support), the vector \mathbf{z}^m is then normally distributed with zero mean and covariance matrix $\mathbf{D}^m (\mathbf{D}^m)^T = \mathbf{I}_m$. In this case, we can compute exactly the PFA as

$$\begin{aligned} \alpha_m &= 1 - \Pr(\max \mathbf{z}^m \leq \eta) = 1 - \Pr(z_1^m \leq \eta)^m, \\ &= 1 - \Phi(\eta)^m, \end{aligned} \quad (14)$$

where Φ is the cumulative function of the normal distribution. In practice, the dictionary is chosen to be highly coherent (as we want to track close translated versions of a reference spectrum). This requires another way to be found to estimate or maximize this probability.

Proposition III.1. For any $t \in \mathbb{R}$ and $m \geq 2$, let $M_{m+1}(t)$ be defined recursively, under \mathcal{H}_0 , as

$$M_{m+1}(t) = \Pr(z_1^{m+1} \leq t \mid z_2^{m+1} \leq t, z_3^{m+1} \leq t) \times M_m(t), \quad (15)$$

with $M_2(t) = \Pr(z_1^2 \leq t, z_2^2 \leq t)$. Under the aforementioned assumptions, an upper boundary of the PFA α_m is given by $1 - M_m(\eta)$.

Proof. See Appendix C. \square

The interest of expression (15) is that the first factor of the right hand side and the initial value $M_2(t)$ can be evaluated numerically based on quadrature rules for trivariate and bivariate normal distribution functions [30], without the requirement for any Monte-Carlo approximation. Thus, this upper boundary can be easily and accurately computed. When the atoms are uncorrelated, this boundary is sharp and reduces to (14). Moreover this allows appreciation of, for a given threshold η , the increase in the PFA α_m as a function of the dictionary size m for highly correlated atoms. In a reciprocal way, this allows the evaluation of the threshold η_m , which ensures a false alarm rate that is lower than a given α for any $m \geq 1$.

We can now estimate roughly the potential detection gain under \mathcal{H}_1 as a function of m , as follows: Under \mathcal{H}_1 , we have assumed that

$$\mathbf{y} = a\mathbf{d}_*^u + \epsilon$$

with a shift $u \sim \mathcal{U}([-\tau, \tau])$. Then, if we assume that the maximum is obtained for the closest atom, which is by assumption the more correlated with \mathbf{d}_*^u , the expected max-test statistic can be approximated by

$$E[\max \mathbf{z}^m] \approx aE[\Gamma(e_m)]$$

where $\Gamma(\cdot)$ is the autocorrelation function of \mathbf{d}_* and $e_m \sim \mathcal{U}([0, \tau/(m-1)])$ is the shift between \mathbf{d}_*^u and the closest atom.

Using this expected max-test value under \mathcal{H}_1 and the upper boundary on the false alarm given in proposition III.1, we can see in Figure 8 that when the dictionary size m increases, the max-test statistic can still increase under \mathcal{H}_1 . However, for fixed level control α , the test threshold (upper boundary) η_m does not increase significantly above a certain size, e.g., $m \geq 10$ in Figure 8. This is clearly explained by the stronger correlations when adding new atoms. Conversely, if the atoms are uncorrelated, we see that the threshold deduced from (14) increases faster than the potential gain of the max-test statistic under \mathcal{H}_1 .

This is confirmed in Figure 7, which shows different empirical ROC curves for different sized m of the LSS dictionaries. It can be seen empirically that the more coherent the dictionary, the more powerful becomes the max test.

As a consequence, in the application, the dictionary will be built to be as coherent as possible for the spectral resolution of the MUSE instrument. The reference atom \mathbf{d}_* is estimated by averaging the spectra of the 5 pixels at the spatial intensity peak of the galaxy. The spectrum is limited to a $l = 30$ spectral band area centered on the spectral emission peak, which ensures the presence of the whole emission line feature. Based on astronomical priors, the spectral shift is limited to the interval $[-\tau, \tau]$ with $\tau = 7$ MUSE spectral bands (i.e., $\tau \approx 9 \text{ \AA}$). Shifting is done at the spectral resolution of the instrument, to avoid any interpolations. The dictionary \mathbf{D}^m is finally built with the atoms corresponding to these $m = 15$ shifted versions of \mathbf{d}_* .




Fig. 7: Comparison of empirical ROC curves for several dictionary sizes m under the LSS dictionary model. The results were obtained on 50 Monte-Carlo runs of simulated MUSE-like data under the aforementioned assumptions.




Fig. 8: Test statistics threshold versus the number of dictionary atoms m for a control level of $\alpha = 0.05$. \blacklozenge marker curve, threshold η_m for the LSS dictionary \mathbf{D}^m as in Fig. 6; \bullet marker curve, threshold for a size m dictionary with uncorrelated atoms; heavy blue curve, evolution of the test statistic potential gain under \mathcal{H}_1 for an intensity $a = 2.7$.

2) *Similarity measure:* To test whether a given spectrum belongs to the extended source, the similarity measure used in this application is the SAD, as defined in (5). Of note, other metrics were explored to build the test statistic, such as the matched filter one (4) and the spectral information divergence defined in [31]. Spectral information divergence is built upon the symmetrical Kullback-Leibler divergence, and it compares the spectra as distribution densities. As it demands positive signals for its computation, it cannot be used directly for our problem, as the MUSE data can be negative due to high symmetrical noise levels. The matched filter approach can be used on the MUSE data, and it gives good results. However SAD appears to be more robust to some systematics of the MUSE data cubes, such as the edges where there is higher variability, and it is preferred here.

D. Results on real data

The results on several subcubes of $n = 50 \times 50$ pixels by $l = 30$ wavelengths centered around interesting objects of the HDFS catalog are shown in Figure 9 for the detection procedure described in Alg. 1. For each of the $n = 50 \times 50 = 2500$ spectra, the max-test statistics (6) are obtained from the SAD similarity measure, and with a highly coherent dictionary constructed as described in the last paragraph of section III-C1. The empirical null estimators are computed on larger subcubes (centered around the subcube to be tested) that are composed of 200 by 200 spectra. For each object, the first row shows the narrow band image around the emission line (the data subcube is totalled along the $l = 30$ spectral bands centered on the emission line peak). The second column of the first row shows the same narrow-band image, but after the different preprocessing steps (which include continuum subtraction and FSF convolution). The last column shows the reference spectrum d_* , built from the pixels at the center of the studied object. On the second row, the first column shows the maps of the empirical p -values (12) obtained for the $n = 2500$ spectra. The maps of the q -values are depicted on the second column. Q -values were introduced in [32] and can be seen as the FDR counterpart of the p -values. For each test statistic, it is defined as the minimal FDR that allows this test statistic to be a discovery. In our detection framework, this is the minimal FDR control level $q \in [0, 1]$ in Alg. 1, such that a given spectrum is detected as part of the halo. The interest in this global measure of significance is clear here, as it allows us to present more contrasted significance maps than the classical p -value maps. The third column shows the binary detection map provided by Alg. 1 for a nominal FDR control level $q = 0.2$. The contours of the detection region for different values of q are also superimposed. From these maps, it can be seen that several of these objects show clear asymmetry, and that they extend beyond the simple support of a punctual source (the black circles show the support of an estimated FSF). Studies are currently being conducted by astronomers at the Centre de Recherche Astrophysique de Lyon to analyze these results and to apply the method to other sky fields.

IV. CONCLUSIONS

In this paper, a new method is proposed to answer a detection problem of a weak target signature that is partially known, but with a possible large variability within an unknown background that is difficult to model. To answer to this problem, an unsupervised detector was proposed, based on a maximum test approach, as studied in [13]. This detector takes explicitly the possible target variability into account by using a highly coherent dictionary. It does not need any knowledge of the background, but a simple noise symmetry assumption, and the non-negativity of the sparse representation of the targeted signal. This allows to estimate the test statistic distribution and to implement a simple detection procedure robust to model/background miss-specification. Moreover, the error control was developed based on a false discovery rate approach, and a global measure of the significance was obtained. Such a control with detection threshold that adapts to the data

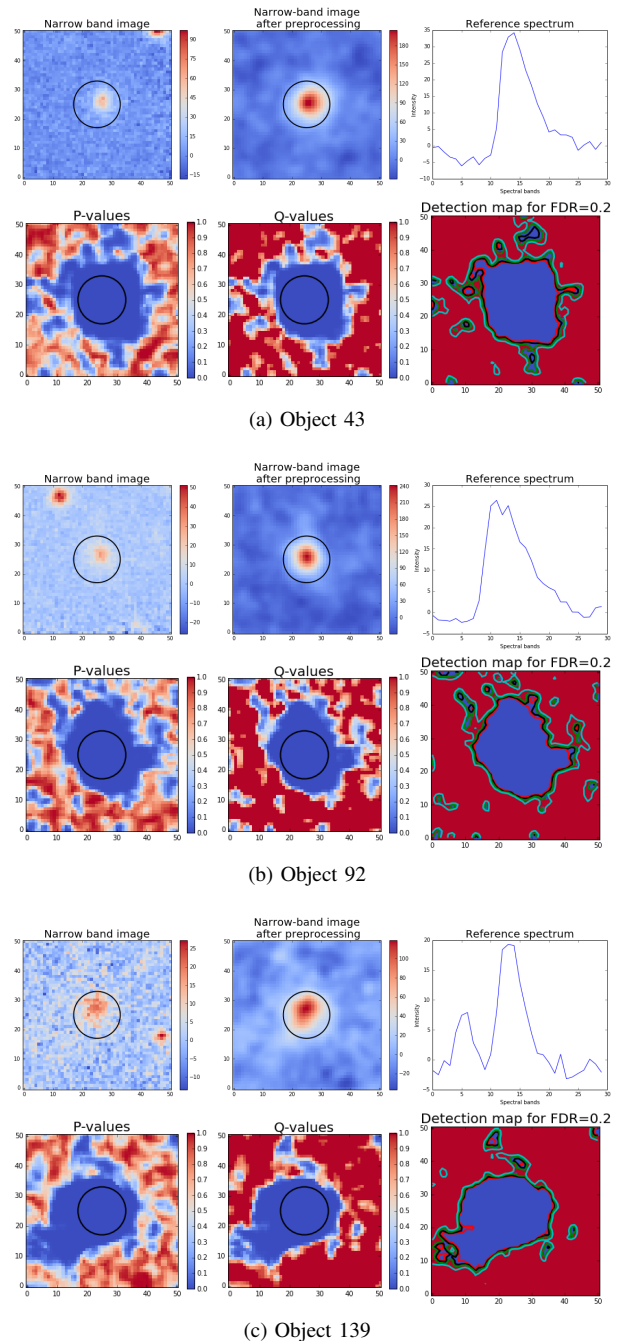


Fig. 9: Results on several objects of the HDFS. Each subfigure corresponds to an object with the numbering to the catalog defined in [27]. For each subfigure: on the first row, from left to right: narrow-band image, narrow-band image after preprocessing (continuum subtraction and FSF convolution), and reference spectrum atom. On the second row, from left to right: p -value map resulting from the test, q -value map, and detection contours for several FDR levels q (0.05 in red, 0.1 in black, 0.2 in green, and 0.4 in cyan) superimposed on the binary detection map for a FDR level $q = 0.2$ (blue pixels, halo detections). The circles in black in the first two columns show the extent of the FSF. Results are analyzed in III-D.

is not yet widely used in the signal processing community whereas it is highly pertinent for processing massive datasets. This whole new process was tested on real MUSE data. The promising obtained results are presently analyzed by astronomers. Future extensions of this original method will account for the existence of spatial structure of the target while controlling the FDR. The MUSE data used in this paper is now publicly available at <http://muse-vlt.eu/science/hdfs-v1-0/>, and the Python code of the proposed method is available on demand.

ACKNOWLEDGMENTS

The authors would like to thank the ERC 339659-MUSICOS for its funding, Roland Bacon and Floriane Leclercq for their expertise on the MUSE data, and Jean-Baptiste Courbot for numerous interesting exchanges on weak target detection.

APPENDIX

Appendix A: Proof for the empirical null estimators

Appendix A1: Proof of lemma II.2: Let $\mathbf{g} = (T_{\max}(\mathbf{y}_1), \dots, T_{\max}(\mathbf{y}_n))$ be the set composed of the n max statistics, the elements of which are denoted as g_i for $1 \leq i \leq n$. Similarly $\mathbf{s} = (-T_{\min}(\mathbf{y}_1), \dots, -T_{\min}(\mathbf{y}_n))$ is the set composed of the n opposite min statistics, the elements of which are denoted as s_i for $1 \leq i \leq n$.

We first show that $\hat{\mu}_0$ verifies (8), i.e., that $\#\{g_i \leq \hat{\mu}_0\} = \#\{s_i > \hat{\mu}_0\}$. For absolutely continuous distributions, $\Pr(t_{(n)} = t_{(n+1)}) = 0$. Thus from (9), we get that $\#\{t_i \leq \hat{\mu}_0\} = n$ with probability one. The sample set \mathbf{t} is the union of \mathbf{g} and \mathbf{s} : if $m_0 = \#\{g_i \leq \hat{\mu}_0\}$, then $\#\{s_i \leq \hat{\mu}_0\} = n - m_0$. As a consequence, $\#\{s_i > \hat{\mu}_0\} = n - (n - m_0) = m_0 = \#\{g_i \leq \hat{\mu}_0\}$, which shows that $\hat{\mu}_0$ verifies (8).

We show now that $\hat{\mu}_0$ converges in probability toward μ_0 . As $\hat{\mu}_0$ satisfies (8), and $\bar{F}(t)$ (resp. $\bar{G}(t)$) converges in probability to $F(t)$ (resp. $G(t)$) for any $t \in \mathbb{R}$, $\hat{\mu}_0$ converges in probability to the solution of

$$F(t) = G(t).$$

if this equation admits a unique solution. Assuming that the median of F_0 is uniquely defined, it follows that for $t > \mu_0$

$$F(t) = \pi_0 F_0(t) + \pi_1 F_1(t) \geq \pi_0 F_0(t) > \pi_0 F_0(\mu_0) = \pi_0/2,$$

Moreover, for $t > \mu_0$

$$G(t) = \pi_0 G_0(t) < \pi_0 G_0(\mu_0) = \pi_0/2 < F(t),$$

where the first equality is due to zero assumption A4. As a consequence, there is no solution of $(\mu_0, +\infty)$. Similarly, according to zero assumption A3, that for $t < \mu_0$, $F(t) < G(t)$. Therefore the unique solution is for $t = \mu_0$, where $F(\mu_0) = \pi_0 F_0(\mu_0) = \pi_0/2 = \pi_0 G_0(\mu_0) = G(\mu_0)$, which concludes the proof. \square

Appendix A2: Proof of proposition II.3: We first show that the π_0 estimator given in (10) is consistent. From lemma II.2, $\hat{\mu}_0$ converges in probability toward μ_0 : $\hat{\mu}_0 \xrightarrow{P} \mu_0$. The triangular inequality ensures that $|\bar{F}(\hat{\mu}_0) - F(\mu_0)| \leq |\bar{F}(\hat{\mu}_0) - F(\hat{\mu}_0)| + |F(\hat{\mu}_0) - F(\mu_0)|$. The first term of the right-hand side is dominated by $\sup_t |\bar{F}(t) - F(t)|$, which converges in probability toward 0, according to assumption A5. The second term also converges in probability toward 0, according to the continuous mapping theorem. Thus $\bar{F}(\hat{\mu}_0) \xrightarrow{P} F(\mu_0)$.

According to (7), $F(\mu_0) = \pi_0 F_0(\mu_0) = \frac{\pi_0}{2}$. As $2\bar{F}(\hat{\mu}_0) = 2\frac{n_0}{n}$, this shows that

$$\tilde{\pi}_0 \equiv 2\frac{n_0}{n} \xrightarrow{P} \pi_0 \in (0, 1].$$

Thus $\hat{\pi}_0 = \min\{\tilde{\pi}_0, 1\}$ also converges in probability to π_0 .

We show now the consistency of (11) for $t \in \mathbb{R}$. From (7) and assumption A5, it follows now that $\bar{F}(t) \xrightarrow{P} \pi_0 F_0(t)$ for all $t \leq \mu_0$. Then, according to the Slutsky theorem, $\bar{F}(t)/\tilde{\pi}_0 \xrightarrow{P} F_0(t)$. As $\hat{F}_0(t) = \bar{F}(t)/\tilde{\pi}_0$ for all $t \leq \mu_0$, this shows the consistency for $t \leq \mu_0$. The demonstration for $t > \mu_0$ can be done in a similar manner, by noting that $\hat{F}_0(t) = 1 - \bar{G}(t)/\tilde{\pi}_0$ for $t > \mu_0$. \square

Appendix B: Proof of proposition II.4

Let $T_{\max}(\mathbf{y}_{(1)}) \leq T_{\max}(\mathbf{y}_{(2)}) \leq \dots \leq T_{\max}(\mathbf{y}_{(n)})$ be the ordered max-test statistics, while $p_{(1)} \leq p_{(2)} \leq \dots \leq p_{(n)}$ are the ordered p -values (while p_i notes the p -value associated with pixel i). From (12), it follows that,

$$p_{(i)} = 1 - \hat{F}_0\left(T_{\max}(\mathbf{y}_{(j)})\right) \quad \text{for } 1 \leq i \leq n$$

where $j = n - i + 1$. For $j \leq n_0$, $T_{\max}(\mathbf{y}_{(j)}) \leq \mu_0$ thus $\#\{s_{0,i} \leq T_{\max}(\mathbf{y}_{(j)})\} = j$ and $\#\{g_{0,i} \leq T_{\max}(\mathbf{y}_{(j)})\} = 0$. So, $2n_0 \hat{F}_0\left(T_{\max}(\mathbf{y}_{(j)})\right) = j$ for $j \leq n_0$, that is $2n_0 p_{(i)} = 2n_0 - n + i - 1$ for $i \geq n - n_0 + 1$. For $k \geq n_0$, let $i_k = n - 2n_0 + 1 + k$. Then $i_k \geq n - n_0 + 1$ so $2n_0 p_{(i_k)} = 2n_0 - n + i_k - 1 = k$. Thus $\#\{2n_0 p_i > k\} = n - i_k = 2n_0 - k - 1$. If $\zeta = k/2n_0$, then $\#\{p_i > \zeta\} = \#\{2n_0 p_i > k\}$. Thus $\frac{1 + \#\{p_i > \zeta\}}{(1-\zeta)n} = \frac{2n_0 - k}{(1-k/2n_0)n} = \frac{2n_0}{n}$, which shows that $\hat{\pi}_0^*(\zeta) = \hat{\pi}_0$.

In the general case where $\zeta \in [\frac{1}{2}, 1)$, then ζ and $k_\zeta/2n_0$, where $k_\zeta = \lfloor 2n_0 \zeta \rfloor \in \{n_0, \dots, 2n_0 - 1\}$, are asymptotically equivalent when n_0 grows to infinity. Thus, $\hat{\pi}_0^*(\zeta)$ and $\hat{\pi}_0^*(k_\zeta/2n_0) = \hat{\pi}_0$ are asymptotically equivalent. This concludes the proof. \square

Appendix C: Proof of proposition III.1

Under \mathcal{H}_0 , for a threshold t we have:

$$\begin{aligned} \Pr(\max \mathbf{z}^{m+1} \leq t) &= \Pr(z_1^{m+1} \leq t, \dots, z_{m+1}^{m+1} \leq t) \\ &= \Pr(z_1^{m+1} \leq t \mid z_2^{m+1} \leq t, \dots, z_{m+1}^{m+1} \leq t) \\ &\quad \times \Pr(z_2^{m+1} \leq t, \dots, z_{m+1}^{m+1} \leq t) \end{aligned}$$

As $\mathbf{D}^m \geq 0$ and $\mathbf{y} \sim \mathcal{N}(0, \mathbf{I}_m)$ under \mathcal{H}_0 , \mathbf{z}^m is positively associated in the sense of [33]. Thus,

$$\begin{aligned} \Pr(z_1^{m+1} \leq t \mid z_2^{m+1} \leq t, \dots, z_{m+1}^{m+1} \leq t) &\geq \\ \Pr(z_1^{m+1} \leq t \mid z_2^{m+1} \leq t, z_3^{m+1} \leq t) & \quad (16) \end{aligned}$$

Using the numerical procedures given in [30], we can accurately compute the right-hand side term of (16). Note that this term gives a relatively sharp lower boundary, because z_2^{m+1} and z_3^{m+1} are the more correlated variables with z_1^{m+1} among the z_j^{m+1} for $j \geq 2$.

Moreover, by construction, the shifts between the atoms in D^{m+1} are smaller than those for the atoms in D^m . With the autocorrelation function assumed to be non-increasing with the absolute shifts, it follows that the size m Gaussian random vector $(z_2^{m+1}, \dots, z_{m+1}^{m+1})$ has larger correlations than the size m Gaussian random vector (z_1^m, \dots, z_m^m) . By assumption, these two vectors are centered with unit marginal variances under \mathcal{H}_0 . Thus the Slepian lemma [34] yields that:

$$\Pr(z_2^{m+1} \leq t, \dots, z_{m+1}^{m+1} \leq t) \geq \Pr(z_1^m \leq t, \dots, z_m^m \leq t). \quad (17)$$

By combining (16) and (17), we can then minimise $\Pr(\max z^m \leq t)$ by a function $M_m(t)$ that is defined recursively as:

$$M_{m+1}(t) = \Pr(z_1^{m+1} \leq t \mid z_2^{m+1} \leq t, z_3^{m+1} \leq t) \times M_m(t),$$

where $M_2(t) = \Pr(z_1^2 \leq t, z_2^2 \leq t)$. This gives the upper boundary for the PFA of proposition III.1. Numerical computations emphasize that $M_m(t)$ increases with t . Then it is possible to (numerically) inverse $M_m(\eta)$ to get η_m for a control level α : $\eta_m = M_m^{-1}(1 - \alpha)$ verifies $\Pr(\max z^m > \eta_m) \leq \alpha$. \square

REFERENCES

- [1] R. Bacon, M. Accardo, L. Adjali, H. Anwand, S. Bauer, I. Biswas, J. Blaizot, D. Boudon, S. Brau-Nogue, J. Brinchmann *et al.*, "The muse second-generation vlt instrument," in *SPIE Astronomical Telescopes+ Instrumentation*. International Society for Optics and Photonics, 2010, pp. 773 508–773 508.
- [2] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 289–300, 1995.
- [3] D. Manolakis, R. Lockwood, T. Cooley, and J. Jacobson, "Is there a best hyperspectral detection algorithm?" in *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics, 2009, pp. 733 402–733 402.
- [4] I. S. Reed and X. Yu, "Adaptive multiple-band cfar detection of an optical pattern with unknown spectral distribution," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 38, no. 10, pp. 1760–1770, 1990.
- [5] D. G. Manolakis, G. A. Shaw, and N. Keshava, "Comparative analysis of hyperspectral adaptive matched filter detectors," in *AeroSense 2000*. International Society for Optics and Photonics, 2000, pp. 2–17.
- [6] L. L. Scharf and L. T. McWhorter, "Adaptive matched subspace detectors and adaptive coherence estimators," in *Signals, Systems and Computers, 1996. Conference Record of the Thirtieth Asilomar Conference on*. IEEE, 1996, pp. 1114–1117.
- [7] J. Solomon and B. Rock, "Imaging spectrometry for earth remote sensing," *Science*, vol. 228, no. 4704, pp. 1147–1152, 1985.
- [8] Y. Chen, N. M. Nasrabadi, and T. D. Tran, "Sparse representation for target detection in hyperspectral imagery," *IEEE Journal of Selected Topics in Signal Processing*, vol. 5, no. 3, pp. 629–640, 2011.
- [9] S. Paris, D. Mary, and A. Ferrari, "Detection tests using sparse models, with application to hyperspectral data," *Signal Processing, IEEE Transactions on*, vol. 61, no. 6, pp. 1481–1494, 2013.
- [10] J.-B. Courbot, V. Mazet, E. Monfrini, and C. Collet, "Detection of faint extended sources in hyperspectral data and application to hdf-s muse observations," in *IEEE International Conference on Acoustic, Speech and Signal Processing 2016*, 2016.
- [11] D. Donoho and J. Jin, "Higher criticism for detecting sparse heterogeneous mixtures," *The Annals of Statistics*, vol. 32, no. 3, pp. 962–994, 2004.
- [12] P. Hall, J. Jin *et al.*, "Innovated higher criticism for detecting sparse signals in correlated noise," *The Annals of Statistics*, vol. 38, no. 3, pp. 1686–1732, 2010.
- [13] E. Arias-Castro, E. J. Candès, and Y. Plan, "Global testing under sparse alternatives: Anova, multiple comparisons and the higher criticism," *The Annals of Statistics*, pp. 2533–2556, 2011.
- [14] D. Donoho and J. Jin, "Higher criticism for large-scale inference, especially for rare and weak effects," *Statistical Science*, vol. 30, no. 1, pp. 1–25, 2015.
- [15] P. Serra, R. Jurek, and L. Flöer, "Using negative detections to estimate source-finder reliability," *Publications of the Astronomical Society of Australia*, vol. 29, no. 3, pp. 296–300, 2012.
- [16] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [17] K. Huang and S. Aviyente, "Sparse representation for signal classification," in *Advances in Neural Information Processing Systems 19: Proceedings of the NIPS 2006 Conference*. MIT Press, 2007, pp. 609–616.
- [18] R. A. Schowengerdt, *Remote sensing: models and methods for image processing*. Academic press, 2006.
- [19] C. Meillier, F. Chatelain, O. Michel, and H. Ayasso, "Nonparametric bayesian extraction of object configurations in massive data," *IEEE Transactions on Signal Processing*, vol. 63, no. 8, pp. 1911–1924.
- [20] C. Meillier, "Détection de sources quasi-ponctuelles dans des champs de données massifs," Ph.D. dissertation, Université Grenoble Alpes, 2015.
- [21] B. Efron, *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*. Cambridge University Press, 2012, vol. 1.
- [22] C. Meillier, F. Chatelain, O. Michel, and H. Ayasso, "Error control for the detection of rare and weak signatures in massive data," in *Signal Processing Conference (EUSIPCO), 2015 23rd European*. IEEE, 2015, pp. 1974–1978.
- [23] C. R. Genovese, N. A. Lazar, and T. Nichols, "Thresholding of statistical maps in functional neuroimaging using the false discovery rate," *Neuroimage*, vol. 15, no. 4, pp. 870–878, 2002.
- [24] Y. Benjamini and D. Yekutieli, "The control of the false discovery rate in multiple testing under dependency," *Annals of statistics*, pp. 1165–1188, 2001.
- [25] J. D. Storey, J. E. Taylor, and D. Siegmund, "Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 66, no. 1, pp. 187–205, 2004.
- [26] C. Meillier, F. Chatelain, O. Michel, R. Bacon, L. Piqueras, R. Bacher, and H. Ayasso, "Selfi: an object-based, bayesian method for faint emission line source detection in muse deep field data cubes," *A& A*, vol. 588, p. A140, 2016. [Online]. Available: <http://dx.doi.org/10.1051/0004-6361/201527724>
- [27] R. Bacon, J. Brinchmann, J. Richard, T. Contini, A. Drake, M. Franx, S. Tacchella, J. Vernet, L. Wisotzki, J. Blaizot *et al.*, "The MUSE 3D view of the Hubble Deep Field South," *A&A*, vol. 575, p. A75, 2015.
- [28] R. Bacher, F. Chatelain, and O. Michel, "An adaptive robust regression method: application to galaxy spectrum baseline estimation," in *IEEE International Conference on Acoustic, Speech and Signal Processing 2016*, 2016.
- [29] E. Villeneuve, H. Carfantan, and D. Serre, "PSF estimation of hyperspectral data acquisition system for ground-based astrophysical observations," in *2011 3rd Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*. IEEE, 2011, pp. 1–4.
- [30] A. Genz and F. Bretz, *Computation of multivariate normal and t probabilities*. Springer Science & Business Media, 2009, vol. 195.
- [31] C.-I. Chang, "Spectral information divergence for hyperspectral image analysis," in *Geoscience and Remote Sensing Symposium, 1999. IGARSS'99 Proceedings. IEEE 1999 International*, vol. 1. IEEE, 1999, pp. 509–511.
- [32] J. D. Storey and R. Tibshirani, "Statistical significance for genomewide studies," *Proceedings of the National Academy of Sciences*, vol. 100, no. 16, pp. 9440–9445, 2003.
- [33] J. D. Esary, F. Proschan, D. W. Walkup *et al.*, "Association of random variables, with applications," *The Annals of Mathematical Statistics*, vol. 38, no. 5, pp. 1466–1474, 1967.
- [34] D. Slepian, "The one-sided barrier problem for gaussian noise," *Bell System Technical Journal*, vol. 41, no. 2, pp. 463–501, 1962.