
Quantized nonnegative matrix factorization

de Fréin, Ruairí

Amadeus SAS, Sophia Antipolis
Nice, France

&

KTH Royal Institute of Technology, Sweden
web: <https://robustandscaleable.wordpress.com>

in: Digital Signal Processing (DSP), 2014 19th IEEE Int. Conf.. See also $\text{BIB}_{\text{T}}\text{E}_\text{X}$ entry below.

$\text{BIB}_{\text{T}}\text{E}_\text{X}$:

```
@article{deFrein14Quantized,  
author={de Fr\`{e}in, Ruair\`{i}},  
booktitle={Digital Signal Processing (DSP), 2014 19th IEEE Int. Conf.},  
title={Quantized nonnegative matrix factorization},  
year={2014},  
pages={377--382},  
keywords={encoding;iterative methods;learning (artificial intelligence);  
matrix decomposition;quantisation (signal);Frobenius-norm quantized nonnegative  
matrix factorization algorithm;QNMf;adaptive quantization levels;decomposition  
rank of interest;element-wise quantization constraints;encoding techniques;  
extended NMF iteration;factor signal values set;learning algorithm;  
matrix decomposition task;outer quantization optimization;  
post factorization quantization;quantization residual;quasifixed quantization levels;  
rank reduction;signal compaction;signal-to-noise-ratios;Approximation methods;  
Convergence;Digital signal processing;Matrix decomposition;Optimization;  
Quantization (signal);Signal processing algorithms;low rank;nmf;quantization},  
doi={10.1109/ICDSP.2014.6900690},  
url={URL: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6900690&isnumber=6900665},  
month={Aug}, }
```

© 2014 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.



Quantized Nonnegative Matrix Factorization

Ruairí de Fréin

Telecommunications Software and Systems Group
Ireland

Email: rdefrein@gmail.com

Abstract—Even though Nonnegative Matrix Factorization (NMF) in its original form performs rank reduction and signal compaction implicitly, it does not explicitly consider storage or transmission constraints. We propose a Frobenius-norm Quantized Nonnegative Matrix Factorization algorithm that is 1) almost as precise as traditional NMF for decomposition ranks of interest (with in 1-4dB), 2) admits to practical encoding techniques by learning a factorization which is *simpler* than NMF's (by a factor of 20-70) and 3) exhibits a complexity which is comparable with state-of-the-art NMF methods. These properties are achieved by considering the quantization residual via an outer quantization optimization step, in an extended NMF iteration, namely QNMF. This approach comes in two forms: QNMF with 1) quasi-fixed and 2) adaptive quantization levels. Quantized NMF considers element-wise quantization constraints in the learning algorithm to eliminate defects due to post factorization quantization. We demonstrate significant reduction in the cardinality of the factor signal values set for comparable Signal-to-Noise-Ratios in a matrix decomposition task.

Index Terms—low rank; nmf; quantization;

I. INTRODUCTION

Nonnegative Matrix Factorization (NMF) is a fundamental tool in Signal Processing and Machine Learning, which is used for portfolio optimization [1] and Blind Source Separation [2], [3], [4]. The appeal of NMF is that it learns an adaptive *basis* (set of vectors) and sparse coefficients: NMF represents an input stimulus ensemble as a linear combination of elements from a representative set of learned NMF *basis* functions.

Given the matrix \mathbf{V} , NMF decomposes \mathbf{V} into the product of two matrices, $\mathbf{W} \in \mathfrak{R}_+^{M \times R}$ and $\mathbf{H} \in \mathfrak{R}_+^{R \times N}$. All matrices have exclusively nonnegative elements, $M > R, N > R$. NMF-Frobenius' objective is the squared- ℓ_2 norm:

$$f(\mathbf{V} \|\mathbf{WH}) = \frac{1}{2} \sum_{m,n} |\mathbf{V}_{m,n} - (\mathbf{WH})_{m,n}|^2. \quad (1)$$

A suitable step-size parameter, proposed by Lee and Seung in [5], results in two alternating, multiplicative, gradient descent updating algorithms:

$$\mathbf{W} \leftarrow \mathbf{W} \odot \mathbf{VH}^T \oslash \mathbf{WHH}^T, \quad (2)$$

$$\mathbf{H} \leftarrow \mathbf{H} \odot \mathbf{W}^T \mathbf{V} \oslash \mathbf{W}^T \mathbf{WH} \quad (3)$$

where \odot represents element-wise multiplication, and \oslash is element-wise division. The NMF solution is generally not unique [6], [7], or exact. For every invertible \mathbf{A} we have a potential factorization [8], [9],

$$\mathbf{V} \approx (\mathbf{WA})(\mathbf{A}^{-1}\mathbf{H}). \quad (4)$$




Fig. 1. QNMF reconstruction: The input matrix \mathbf{V} is real valued and factorized into two real-valued factors \mathbf{W}, \mathbf{H} in the NMF block. The NMF factors are then quantized using either fixed step-size quantization or adaptive stepsize quantization producing integer valued factors $\hat{\mathbf{W}} \in \mathcal{Z}^{M \times R}, \hat{\mathbf{H}} \in \mathcal{Z}^{R \times N}$. Adaptive quantization optimization is denoted by the double ended arrows from the step-size to the NMF optimization routine. The resulting factorization has the same dimensions as \mathbf{V} but it is integer valued, $\hat{\mathbf{V}} \in \mathcal{Z}^{M \times N}$

In recent work, [7] the authors *solved* the *uniqueness* problem, and addressed the *inexactness* of NMF, by introducing the idea that rank-1 NMF approximations are closures. Using this restriction, the solution space was constrained to the extent that the approximation could be replaced by equality in (Eqn. 4). The Nonnegative Matrix Approximation (NMA) in (Eqn. 4), described by Lee and Seung in [5] is used in many applications. Typically any member of the set of NMA's that meets an arbitrary accuracy constraint suffices: this apparent *freedom* motivates the present contribution, namely QNMF, which produces factors which have a small set of possible entry values, and are thus amenable to efficient storage (via a supplementary compression process).

Definition 1: A valid NMF solution is a pair of matrices $\{\mathbf{W}, \mathbf{H}\}$ which are element-wise nonnegative $\mathbf{W}_{m,r} \geq 0, \mathbf{H}_{r,n} \geq 0$, that satisfy the accuracy condition $\|\mathbf{V} - \mathbf{WH}\|^2 \leq \eta$, where $\eta \in \mathfrak{R}$ is an arbitrary small positive number. The set of valid NMFs is denoted

$$S = \{\{\mathbf{W}, \mathbf{H}\} \mid \|\mathbf{V} - \mathbf{WH}\|^2 \leq \eta\}. \quad (5)$$

We consider the problem of transmission and storage of the approximation in (Eqn. 4). We use additional criteria to select members of S with appropriate properties –here the crucial property is that the factors are element-wise integer and that this set of integers has low cardinality.

We consider the factorization-quantization-reconstruction scenario described in Fig.1. Our hypothesis is that traditional NMF may yield an inefficient representation for transmission. The inputs matrix $\mathbf{V} \in \mathfrak{R}^{M \times N}$ is factorized-quantized into the product of two integer matrices $\hat{\mathbf{W}} \in \mathcal{Z}^{M \times R}$ and $\hat{\mathbf{H}} \in \mathcal{Z}^{R \times N}$ and two quantization step-size (scalars or) vectors, $\Delta_{\mathbf{W}}$ and

Δ_h . To capture this idea, we contribute a new NMA mixing model which is expressed as:

$$\hat{\mathbf{V}} = \hat{\mathbf{W}} \text{diag} \Delta_w \text{diag} \Delta_h \hat{\mathbf{H}}, \quad \mathbf{V} \approx \hat{\mathbf{V}}. \quad (6)$$

The operation, $\text{diag} \mathbf{x}$, constructs a square matrix with the vector \mathbf{x} 's elements on the diagonal. The factors produced by an *inner* NMF $\mathbf{W} \in \mathfrak{R}^{M \times R}$, $\mathbf{H} \in \mathfrak{R}^{R \times N}$ (for example à la Lee and Seung in [5]) are scalar, or element-wise quantized –quantized by a quantizer which acts separately on each component of \mathbf{W} and \mathbf{H} – as part of the learning routine.

Definition 2: A Quantized NMF solution is a member of the set S_q

$$S_q = \{ \{ \hat{\mathbf{W}}, \hat{\mathbf{H}} \} \mid \| \mathbf{V} - \hat{\mathbf{W}} \text{diag} \Delta_w \text{diag} \Delta_h \hat{\mathbf{H}} \|^2 \leq \eta \}. \quad (7)$$

The crucial observation in this paper, is that if quantization error minimization is alternated with factorization updates, we can learn members of the set S_q which are as accurate as element-wise-continuous-valued NMF.

We have introduced the QNMF mixing model. In § II we motivate NMF in light of related quantized expansions. We examine the effects of quantization on NMF in § III. We contribute a family of QNMF algorithms in § IV. We perform a numerical evaluation of QNMF in § V and discuss future research directions in § VI.

II. RELATED WORK

The requirement for QNMF arises when many measurements can be taken at a sensing location but only a low-rank approximation is required at a receiver/decision making entity or storage entity. For example, the authors of [10] consider the Matrix Completion problem: they recover and classify Wireless Sensor Network data while minimizing the number of samples that is acquired, processed, and transmitted. In essence a low rank matrix is recovered from a small number of randomly sampled entries. However, quantization is not considered in the learning algorithm. More generally, scaling the algorithms [1], [2], [3], [4] to large scale data may be problematic; what is required is a good quality low-rank compressible approximation. How can we best estimate \mathbf{V} from $\hat{\mathbf{W}}$ and $\hat{\mathbf{H}}$ in these situations? How does the quality of the $\hat{\mathbf{V}}$ depend on the properties of $\hat{\mathbf{W}}$, $\hat{\mathbf{H}}$, Δ_w and Δ_h ? These fundamental questions are addressed here for the first time by introducing QNMF. The effects of QNMF with respect to the appropriate perceptual quality measure in these domains is deferred to future work. Other potential applications for QNMF include matrix decomposition/missing data estimation in wireless sensor networks [11].

Entropy-based distortion measures and bit allocations for compression are often based on the mean-squared error [12], [13], [14]. This motivates our selection of the Frobenius-norm for QNMF. Our goal is to learn the parameters $\{ \hat{\mathbf{W}}, \hat{\mathbf{H}}, \Delta_w, \Delta_h \}$ so that the reconstruction error $f = \| \mathbf{V} - \hat{\mathbf{V}} \|_2^2$ is minimized. The properties of an NMF in the presence of coefficient quantization have not been explored. We contribute algorithms to learn NMF of this form by considering quantization as part of the alternating minimizing low-rank

TABLE I
QNMF ALGORITHM AND ITS ACRONYM

QNMF	Acronym	Quantizer	Adaptive
Round random	'roundrand'	round or \sim round	x
Floor random	'floorrand'	[·] or [·]	x
Round random Adaptive	'roundadapt'	[·] or [·]	✓
Floor random Adaptive	'flooradapt'	round or \sim round	✓

factorization. The effects of coefficient quantization on representations in \mathfrak{R}^N using overcomplete sets of vectors are investigated in [14]. We empirically investigate the distortion as a function of the rank parameter R and the cardinality of the set of values the factorization takes (which we call “simplicity”). This is the first paper that examines the effects of quantization on a nonnegative matrix factorization. If NMF is to scale to big data analysis, efficient storage of the factors is one problem that must be addressed. Uniform quantization is selected in this paper because the authors of [15] demonstrate that it has an output entropy which is asymptotically smaller than that of any other quantizer, independent of the density function or the error criterion. An excellent quantization survey is found in [16]. Finally, given that NMF learns basis functions with different scales, we introduce a step-size parameter for each column (row) of \mathbf{W} (\mathbf{H}) to improve its *simplicity*.

We introduce some notation. $\mathcal{U}(0, 1)$ is used to represent a uniform distribution from the domain $(0, 1)$. \mathcal{U}_{NM} represents a matrix of size $M \times N$ –each element is independently drawn from $\mathcal{U}(0, 1)$. The subscript is frequently omitted. A vector of R ones is denoted by $\mathbf{1}_R$.

III. QUANTIZED NMF FRAMEWORK

We propose a Quantized Nonnegative Matrix Factorization algorithm that is 1) almost as precise as traditional NMF, 2) admits to practical encoding techniques 3) and exhibits a complexity which is comparable with state-of-the-art NMF methods in § IV. These properties are achieved by considering the quantization residual, via an outer quantization optimization step, in an extended NMF iteration. The use of the quantized activations, for example, in the optimization of the W-update reduces the propagation of the quantization error to subsequent NMF iterations. We consider the effects of quantization on NMF here.

QNMF Framework: Table ?? describes the quantized NMF framework underpinning the results in this paper. It is composed of an inner NMF optimization routine (consisting of an unshaded row for the h-update and the w-update) in lines 5 and 8, which is wrapped by an *outer* quantization updating routine (the shaded rows in Table ??) in lines 6,7,9,10 and 12. The algorithm presented in Table ?? is simple: 1) the quantization steps are randomly initialized (less than one), and 2) the quantization steps are fixed for the entire NMF iteration. These short-comings conspire to produce a factorization algorithm which may not converge monotonically –we provide two enhancements which improve convergence in § IV. First we motivate these algorithms by examining quantization error.

Classical Statistical Analysis: To understand why monotonicity is an issue, we consider the affects of the outer routine on the inner routine. A first step is to consider the affect of reconstructing a vector $\mathbf{v} = \mathbf{V}_{:,n}$ when the corresponding activation vector estimate $\mathbf{h} = \mathbf{H}_{:,n}$ has been degraded in some unspecified way. Lines 5-7 in Table ?? provide an example of the introduction of quantization error. Our analysis provides an understanding of the affects of the outer quantization routine where degradation due to quantization is modelled as additive white noise. After each inner NMF h-update (line 5) we have

$$\mathbf{v} = \mathbf{W}\mathbf{h} + \mathbf{E}, \quad (8)$$

where the functions \mathbf{W} are assumed fixed and \mathbf{E} captures element-wise Gaussian noise in the estimate. Suppose we want to approximate \mathbf{v} given $(\mathbf{h} + \beta_h)$ instead of \mathbf{h} (cf. line 7), this approximation is denoted $\hat{\mathbf{v}}$. We assume each $\beta_{r,h}$, where $\beta_h = [\beta_{1,h}, \dots, \beta_{r,h}, \dots, \beta_{R,h}]$, is an independent random variable with zero mean and variance, σ_h^2 . We subtract the estimates to calculate the Mean-Squared-Error

$$(\mathbf{W}\mathbf{h} + \mathbf{E}) - (\mathbf{W}\mathbf{h} + \mathbf{W}\beta_h + \mathbf{E}) = -\mathbf{W}\beta_h. \quad (9)$$

$$\text{MSE}_h = \mathbb{E}\|\mathbf{v} - \hat{\mathbf{v}}\|^2 = \mathbb{E}\|\mathbf{W}\beta_h\|^2 = \mathbb{E}\beta_h^T \mathbf{W}^T \mathbf{W} \beta_h. \quad (10)$$

Re-ordering the summation and the expectation we get

$$\text{MSE}_h = \sigma_h^2 \text{Tr} \mathbf{W}^T \mathbf{W}. \quad (11)$$

Proposition 1: Noise reduction of each NMF update: Let \mathbf{W} and \mathbf{H} be the features and activation matrices generated by the NMF step, and let $\beta_w = [\beta_{1,w}, \dots, \beta_{r,w}, \dots, \beta_{R,w}]$ $\beta_h = [\beta_{1,h}, \dots, \beta_{r,h}, \dots, \beta_{R,h}]$ be zero mean independent random variables with variance σ_w^2 and σ_h^2 respectively. The MSE of the NMF updates is

$$\text{MSE}_h = \sigma_h^2 \text{Tr} \mathbf{W}^T \mathbf{W}, \quad \text{MSE}_w = \sigma_w^2 \text{Tr} \mathbf{H} \mathbf{H}^T.$$

In this paper, degradation is due to scalar quantization:

$$\begin{aligned} \hat{\mathbf{W}}_{m,:} &= Q(\mathbf{W}_{m,:}) \leftarrow [q_1(\mathbf{W}_{m,1}), \dots, q_r(\mathbf{W}_{m,r}), \\ &\quad \dots, q_R(\mathbf{W}_{m,R})]^T \quad Q: \mathbb{R}^{1 \times R} \mapsto \mathcal{Z}^{1 \times R}, \\ &\quad \text{where } q_r: \mathbb{R} \mapsto \mathbb{R}, 1 \leq r \leq R. \\ \hat{\mathbf{H}}_{:,n} &= Q(\mathbf{H}_{:,n}) \leftarrow [q_1(\mathbf{H}_{1,n}), \dots, q_r(\mathbf{H}_{r,n}), \\ &\quad \dots, q_R(\mathbf{H}_{R,n})]^T \quad Q: \mathbb{R}^{R \times 1} \mapsto \mathcal{Z}^{R \times 1}, \\ &\quad \text{where } q_r: \mathbb{R} \mapsto \mathbb{R}, 1 \leq r \leq R. \end{aligned}$$

1.	$\Delta_w = \mathcal{U}_R; \Delta_w = \frac{\Delta_w}{1^T \Delta_w}, 2. \Delta_h = \mathcal{U}_R; \Delta_h = \frac{\Delta_h}{1^T \Delta_h}$
3.	$I = 2000$; -Iteration count
4.	for $1 \leq i \leq I$
5.	$\mathbf{H} = \mathbf{H} \odot \frac{\mathbf{W}^T \mathbf{V}}{\mathbf{W}^T \mathbf{W} \mathbf{H}}$
6.	$\hat{\mathbf{H}} = \lfloor \text{diag}(\frac{1}{\Delta_h} \mathbf{H} + \frac{1}{2}) \rfloor$
7.	$\mathbf{H} = \text{diag}(\Delta_h) \hat{\mathbf{H}} + \epsilon$
8.	$\mathbf{W} = \mathbf{W} \odot \frac{\mathbf{V} \mathbf{H}^T}{\mathbf{W} \mathbf{H} \mathbf{H}^T}$
9.	$\hat{\mathbf{W}} = \lfloor \mathbf{W} \text{diag}(\frac{1}{\Delta_w} + \frac{1}{2}) \rfloor$
10.	$\mathbf{W} = \hat{\mathbf{W}} \text{diag}(\Delta_w) + \epsilon$
11.	$f(\text{iter}) = \ \mathbf{V} - \mathbf{W} \mathbf{H}\ _2$
12.	Outer Quantization Optimization, endfor

We have outlined an analysis of the case where the quantization noise $\hat{\mathbf{h}} - \mathbf{h}$ and $\hat{\mathbf{w}} - \mathbf{w}$ is random, independent in each dimension and uncorrelated with \mathbf{h} and \mathbf{w} respectively. The assumption that quantization error is signal independent, uniformly distributed white noise is not strictly valid [13]. This assumption fails when the amplitude of the signal is comparable to the quantization step-size. QNMF normalizes the step-sizes Δ_w and Δ_h for this reason. QNMF uses a family of quantization functions. It either rounds (or does not round), floors $\lfloor \cdot \rfloor$ or ceils $\lceil \cdot \rceil$ the factors (cf. Table I), e.g.

$$\hat{\mathbf{H}} = \lfloor \text{diag}(\frac{1}{\Delta_h} \mathbf{H} + \frac{1}{2}) \rfloor \quad \hat{\mathbf{W}} = \lfloor \mathbf{W} \text{diag}(\frac{1}{\Delta_w} + \frac{1}{2}) \rfloor;$$

However, the appeal of NMF is that the objective function

$$f(\mathbf{h}) = \arg \min_{\mathbf{h}} \|\mathbf{v} - \mathbf{W}\mathbf{h}^t\|^2 \quad (12)$$

is monotonically decreased by application of the \mathbf{W} or \mathbf{H} update. Perturbation β_h caused by quantization may cause the objective to be increased. In short, if h^{t+1} minimizes $f(\cdot)$, then $f(h^{t+1}) \leq f(h^t)$, but then if application of quantization reverses this ordering,

$$f(Q(h^{t+1})) \leq f(h^t) \text{ or } f(Q(h^{t+1})) > f(h^t), \quad (13)$$

NMF loses its appealing monotonic convergence property. We consider a number of strategies for addressing the problem of quantizing \mathbf{H} and \mathbf{W} such that $f(\mathbf{h})$ and $f(\mathbf{w})$ are minimized. The quantization functions in this paper are uniform mid-tread [16].

IV. MONOTONICALLY DECREASING QNMF

We introduce 1) random-quasi fixed ‘floorrand’ and ‘roundrand’ QNMFs and 2) adaptive QNMFs, namely ‘flooradapt’ and ‘roundadapt’ QNMF (cf. Table. I), for the framework in Table. ?. This constitutes replacing lines 1,2,6, 9 and 12 in Table. ? with a new outer optimization routine. **Random Quasi-Fixed Quantizer:** A probabilistic method for finding the quantization functions that minimize the objective, f , after each NMF update has been applied (in Table. ?) is presented in Table II. This approach is more efficient than exhaustive enumeration of the quantization functions provided that the goal is merely to find an acceptable solution in a fixed amount of time, rather than the best possible solution. This is compatible with the goal of monotonic convergence of NMF. Consider, there are 2^{RN} and 2^{MR} possible quantization functions for \mathbf{H} and \mathbf{W} .

Table. II summarizes the approach for random quasi-fixed ‘floorrand’ quantization of the \mathbf{H} matrix –a similar technique is applied to the \mathbf{W} matrix, but omitted for brevity. We first investigate if a floored quantizer, applied to each element of \mathbf{H} , minimizes f after the NMF \mathbf{H} -update has been applied. If minimization is achieved quantization is applied and QNMF proceeds. If the objective is not minimized by the floored quantization functions we perform Bernoulli trials to generate a new array of random scalar quantizers where the floor and ceil functions are chosen randomly for each element of the quantization function. The probability of choosing a ‘floor’ is

TABLE II
QUASI-FIXED ‘FLOORRAND’ QUANTIZATION

```

0. for  $0 \leq s \leq \infty$ ,  $s$  increasing in steps of 1
1.  Decrease the step-size:  $\Delta_h \leftarrow \Delta_h / (2^s)$ .
2.  for  $1 \geq p \geq 0$  in steps of .1: Bernoulli trials
3.    for  $1 \leq k \leq 100$  in steps of 1
4.      Draw  $RN$  i.i.d. Bernoulli random variables,
5.      the matrix  $\mathbf{X}$  with success probability  $p$ ,
6.      if  $X_{r,n} = 1$  then  $\hat{\mathbf{H}}_{r,n}^{t+1} = \lceil \text{diag} \frac{1}{\Delta_h} \mathbf{H}_{r,n} + \frac{1}{2} \rceil$ 
7.      else  $\hat{\mathbf{H}}_{r,n}^{t+1} = \lfloor \text{diag} \frac{1}{\Delta_h} \mathbf{H}_{r,n} + \frac{1}{2} \rfloor$ 
8.    endif
9.    if  $f(\mathbf{H}^{t+1}) \leq f(\mathbf{H}^t)$ , set  $\mathbf{H} \leftarrow \mathbf{H}^{t+1}$ , break
10.   endif
11.  endfor endfor endfor

```

p ; decreasing this probability increases the probability of introducing a ceil element in the quantization function. Our goal is to slowly migrate away from the original floored quantization function until a suitable quantizer is found. Multiple Bernoulli trials (≈ 100) are used to generate quantization functions with the probability of a floor for each equal to p . If after many random candidate quantization functions have been applied, the objective is not minimized for a fixed quantization step size, Δ_h , the step-size is divided by two and the process is repeated with the smaller step-size. ‘floorrand’ QNMF is described as follows: Lines 6,7,9 and 10 in Table ?? are augmented to include the generation of the new ‘floorrand’ quantizer in Table. II –random quantization is invoked independently for the matrices \mathbf{H} and \mathbf{W} . Once the new step-sizes are generated, they are maintained for the remainder of the optimization unless they are further decreased.

Empirical evidence suggests that fixing $p = 0.5$ and $s = 0$ for the duration of this simulation provides acceptable convergence (cf. the discussion in § V). ‘roundrand’ quasi-fixed QNMF operates in a similar manner: the floor function in Line 6 in Table. II is replaced by the round function. If line 6 uses a round down quantizer in the previous Bernoulli trial, line 7 uses a round-up quantizer with probability p .

Adaptive Quantizer: The quantization step-sizes in the randomized algorithms ‘floorrand’ and ‘roundrand’ QNMF are quasi-fixed –they are only changed if the Bernoulli randomization routine fails to find a valid solution after exhausting all p . We derive new update rules for the step-size parameters for adaptive QNMF using the Frobenius-norm.

$$\|\mathbf{V} - \hat{\mathbf{W}} \text{diag} \Delta_w \text{diag} \Delta_h \hat{\mathbf{H}}\|^2 \quad (14)$$

is non-increasing under the updates rules for the step-size

$$\Delta_{H,r} \leftarrow \Delta_{H,r} \frac{(\mathbf{W}^T \mathbf{V} \odot \hat{\mathbf{H}}) \mathbf{1}_N}{(\mathbf{W}^T \hat{\mathbf{W}} \text{diag} \Delta_H \hat{\mathbf{H}} \odot \hat{\mathbf{H}}) \mathbf{1}_N} \quad (15)$$

$$\Delta_{W,r} \leftarrow \Delta_{W,r} \frac{\mathbf{1}_M (\mathbf{V} \mathbf{W}^T \odot \hat{\mathbf{W}})}{\mathbf{1}_M (\hat{\mathbf{W}} \text{diag} \Delta_W \mathbf{H} \mathbf{H}^T \odot \hat{\mathbf{W}})} \quad (16)$$

Table III lists the ordering of the adaptive QNMF algorithm. Although these updates are guaranteed to monotonically decrease the objective, the quantization step after the initial NMF

TABLE III
ADAPTIVE QUANTIZATION

```

1.  $\mathbf{H}$  update:  $\mathbf{H} \leftarrow \mathbf{H} \odot \mathbf{W}^T \mathbf{V} \odot \mathbf{W}^T \mathbf{W} \mathbf{H}$ 
2a. Quantize:  $\mathbf{H} = \text{diag}(\Delta_h) \mathbf{H}_q + \epsilon$ 
2b. or apply Quasi-fixed Randomized Quantization
3. Apply:  $\Delta_{H,r} \leftarrow \Delta_{H,r} \frac{(\mathbf{W}^T \mathbf{V} \odot \hat{\mathbf{H}}) \mathbf{1}_N}{(\mathbf{W}^T \hat{\mathbf{W}} \text{diag} \Delta_H \hat{\mathbf{H}} \odot \hat{\mathbf{H}}) \mathbf{1}_N}$ 
4. Normalize:  $\Delta_h = \Delta_h \odot (\mathbf{1}_R^T \Delta_h)$ 
5.  $\mathbf{W}$  update:  $\mathbf{W} \leftarrow \mathbf{W} \odot \mathbf{V} \mathbf{H}^T \odot \mathbf{W} \mathbf{H} \mathbf{H}^T$ 
6a. Quantize:  $\mathbf{W} = \mathbf{W}_q \text{diag}(\Delta_w) + \epsilon$ 
6b. or apply Quasi-fixed Randomized Quantization
7. Apply:  $\Delta_{W,r} \leftarrow \Delta_{W,r} \frac{\mathbf{1}_M (\mathbf{V} \mathbf{W}^T \odot \hat{\mathbf{W}})}{\mathbf{1}_M (\hat{\mathbf{W}} \text{diag} \Delta_W \mathbf{H} \mathbf{H}^T \odot \hat{\mathbf{W}})}$ 
8. Normalize:  $\Delta_w = \Delta_w \odot (\mathbf{1}_R^T \Delta_w)$ 

```

\mathbf{H} and \mathbf{W} updates may introduce an error which is larger than the minimization improvement achieved by applying the updates above, and also by NMF. The aggregate effect of quantization and step-size optimization may cause the objective to increase after the application of these operations. We have observed that in practice adaptive QNMF typically converges (for $s = 0$ and fixed). Put simply, if a sufficient number of quantization levels have been assigned to the algorithm QNMF reduces to the original Lee and Seung NMF. In future work we will analyze the relationship between the quantization error and expected error reduction due to the application of an NMF.

The control of algorithm flow of ‘flooradapt’ and ‘roundadapt’ QNMF is summarized in Table III. In lines 2a or b and lines 6a or b either floor or round randomized quantizers are applied to the factors –the quantizer gives its name to the adaptive algorithm: ‘flooradapt’ and ‘roundadapt’ QNMF.

V. EMPIRICAL EVALUATION

Our empirical evaluation explores the efficacy of using QNMF as an algorithm for lossy compression-factorization. We use Donoho and Stodden’s NMF Swimmers database to evaluate QNMF’s convergence. The Swimmers are an element-wise integer-valued data-set $\mathbf{V}_{m,n} \in \{1, 39\}$; NMF does not produce an element-wise integer factorization of this data-set. Taking our inspiration from Occam’s Razor, we seek a simpler factorization –element-wise finite-resolution data-sets call for an NMF algorithm that generates a factorization with a finite number of different values in the entries. Our choice of the Swimmers is motivated by the fact that the correct factorization, and thus rank parameter $R = 16$, is known. Knowledge of R is important in our analysis of QNMF as choosing an R that is too large (small) generates a low(high) factorization distortion. Our goal is to isolate the effect of quantized updates from the rank selection problem, the Swimmers are suitable.

QNMF Convergence and Complexity: We plot the objective function of the variants of QNMF for the first 100 iterations of QNMF in Fig. V when they are applied to the Swimmers data-set with $R = 16$. Lee-Seung’s Frobenius NMF is plotted as a benchmark method. Each factorization is initialized with the same initial matrices. Firstly, the accuracy of all of factorizations, measured using the squared ℓ_2 -norm,

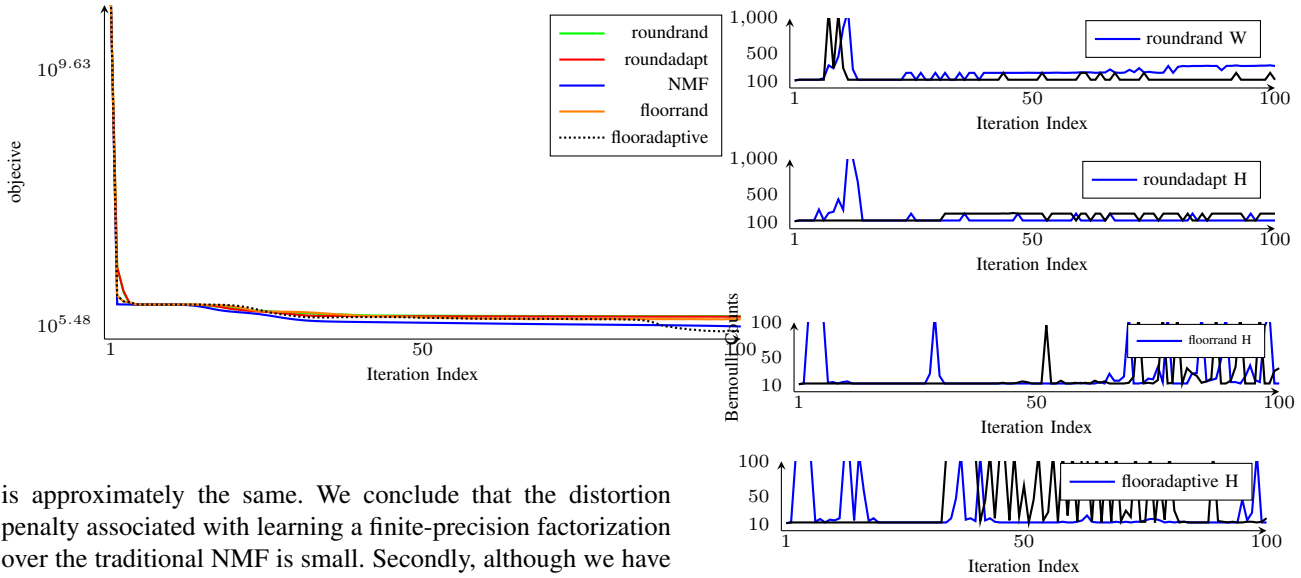


Fig. 2. Bernoulli Counts: number of Bernoulli trials required to minimize the objective (blue \mathbf{H} , black \mathbf{W}).

is approximately the same. We conclude that the distortion penalty associated with learning a finite-precision factorization over the traditional NMF is small. Secondly, although we have not given a proof of convergence for QNMF, the success of these methods in numerous trials (cf. Table IV) suggests that whilst monotonic convergence is not guaranteed the techniques converge to a good *simple* solution. To highlight the convergence property and also the complexity of QNMF, we plot time-series for the factors, \mathbf{W} and \mathbf{H} respectively, which record the counts of the numbers of Bernoulli trials required to decrease the objective function at each iteration. We make the following observations: 1) initially a relatively large number of Bernoulli trials is required for all algorithms; 2) the number of Bernoulli trials reduces to a small number as the quality of the QNMF improves; 3) on first inspection this number of Bernoulli trials may seem onerous, however modern computers have multiple cores and these Bernoulli trials are easily run in parallel; 4) when the Bernoulli count reaches the maximum number of Bernoulli trials (1000 trials for the *rounded* QNMF and 100 trials for *floored* QNMF), the quantization function has failed to minimize the objective. The rounded QNMF algorithms only increase the objective 1 – 3 times. These objective increases occur at the start of the iteration. In some respects this behaviour is analogous to finding a good set of initial matrices, or starting NMF from multiple initial conditions.

The primary factor that causes the objective function to increase is the error introduced by the quantization step. A useful performance measure to evaluate the convergence of QNMF is the Quantization Efficiency, which is defined as: $\omega_Q = SNR_{QNMF}/SNR_{NMF}$ for a given rank. Fig. 3 illustrates the quantization efficiency of the quantization functions computed at each iteration. The quantization efficiency is the ratio of the SNR of the NMF generated updates post and pre quantization. The floored algorithms outperform the rounded algorithms. The convergence in Fig. 3 mirrors that in Fig. V; the variance in the quantization error appears to converge as the distortion measure converges.

QNMF is as precise as NMF for the same rank, but simpler: We demonstrate that QNMF gives a similar quality of decomposition, in terms of the mean SNR, as a NMF




Fig. 3. Quantization Efficiency, ω_Q , of the QNMF quantization functions after each \mathbf{W} and \mathbf{H} update.

of the same rank in Fig. 4 for the Swimmers and for the standard “Cameraman” test image in Table IV. In addition, we illustrate that QNMF learns a simpler decompositions in Fig. 5 for the swimmers and in Table IV for the Cameraman. The input matrix is generated for the Cameraman by blocking (8x8 pixel blocks) the image and placing the vectorized block in the columns of \mathbf{V} . We perform 50 Monte Carlo runs and plot the average SNR and cardinality $C(\{\mathbf{W}\}) + C(\{\mathbf{H}\})$ for a range of ranks $10 \leq R \leq 20$ for the Swimmers and for $10 \leq 40$ for the Cameraman. We compute the cardinality because estimating the differential entropy of NMF or QNMF decompositions, the ideal performance metric, is fraught with difficulty [17]. This discussion underlines the problem with quantizing traditional NMF. Producing a histogram of the




Fig. 4. SNR (dB) of the QNMF algorithms after 100 iterations.




Fig. 5. Count of the number of quantization levels after 100 iterations.

columns (rows) of \mathbf{W} (\mathbf{H}), and then finding the discrete entropy, may produce misleading results because 1) QNMF produces integer factors and NMF produces nonnegative factors, using the same histogram bin selection for these two completely different types of data can generate drastically different discrete entropies; 2) the rows and columns of QNMF and NMF may be permuted and scaled; 3) adaptive QNMFs optimize the quantization step-size across the rows and columns, using a uniform step-size to quantize a NMF to compute its discrete entropy means that we are not comparing like-for-like. In addition, computing the performance of NMF and QNMF with respect to some post-processing compression routine, introduces the possibility that the compression routine favours QNMF. We plot the mean cardinality of the QNMF. The cardinality of the set of entries of NMF is $MR + RN$.

Fig. 4 illustrates that for ranks $10 \leq R \leq 16$ QNMF and NMF give approximately the same accuracy. Table IV shows that QNMF is within 4 dB of NMF. The ‘floorrand’ and ‘flooradapt’ QNMF are the best QNMF algorithms with respect to SNR. For all NMFs the accuracy increases as a function of the rank, over-fitting occurs when $R > 17$ for the Swimmers. The quality of QNMF increases linearly, however NMF experiences a step improvement due to over-fitting. Fig. 5 demonstrates that the adaptive QNMF algorithms exhibit a lower cardinality for approximately the same SNR, in the region $10 \leq R \leq 16$. The factors of these factorizations have approximately 100 fewer quantization levels than the non-adaptive factorizations: ‘roundrand’ and ‘floorrand’. This reduction is significant and implies that they admit to practical encoding techniques. Table IV demonstrates QNMF has a factor of 20-70 fewer levels than NMF for the \mathbf{W} matrix and a factor of 500-1000 fewer levels for \mathbf{H} . Experiments on the Peppers and Barbara images in the supplementary material further support these results.

VI. DISCUSSION

The rationale for the Bernoulli quantization is clear; if the current quantization functions increase the objective, a matrix of quantization functions which have a high probability ($p < 1$) of being a ‘floor’ or ‘round’ but minimize the objective is a satisfactory substitute. Moreover, the number of candidate random quantizers generated for each p is bounded – This bound is much smaller than the total number of possible floor-ceiling matrix permutations, typically $\ll 2^{RN}, 2^{RM}$. In terms of the step-size division process for the round quantizer, when the Δ_h is small (relative to the variation in the signal being

TABLE IV
CAMERAMAN: SNR OF QNMF OF RANKS $10 \leq R \leq 40$. AVERAGE NUMBER OF DIFFERENT VALUES REQUIRED TO BY THE \mathbf{W} AND \mathbf{H} FACTORS. THE NUMBER OF DIFFERENT VALUES USED BY NMF UNDERLINES THE REDUCTION IN CARDINALITY BY QNMF.

Rank R	10	15	20	25	30	35	40
NMF SNR (dB)	24.4	26.6	28.0	29.3	30.0	30.6	31.1
roundrand SNR	18.7	19.9	21.4	22.0	23.2	23.9	24.6
roundadapt SNR	19.4	20.2	21.8	22.6	23.7	24.2	24.8
floorrand SNR	21.3	22.9	24.3	25.2	26.1	26.5	27.2
flooradapt SNR	21.6	22.9	24.2	25.0	25.9	26.2	26.7
NMF $C(\{\mathbf{W}\})$	640	960	1280	1600	1920	2240	2560
roundrand $C(\{\mathbf{W}\})$	37.6	44.7	47.1	49.9	52.1	52.9	53.7
roundadapt $C(\{\mathbf{W}\})$	38.4	43.9	47.3	49.8	51.8	52.5	53.3
floorrand $C(\{\mathbf{W}\})$	31.9	33.0	34.1	35.2	35.6	37.9	39.1
flooradapt $C(\{\mathbf{W}\})$	33.3	33.9	35.4	36.7	37.6	38.9	40.5
NMF $C(\{\mathbf{H}\})$	40960	61440	81920	102400	122880	143360	163840
roundrand $C(\{\mathbf{H}\})$	129.5	135.5	158.8	166.0	197.3	172.6	188.0
roundadapt $C(\{\mathbf{H}\})$	98.6	110.2	129.6	140.5	160.7	156.3	165.4
floorrand $C(\{\mathbf{H}\})$	172.9	187.4	224.4	227.2	257.0	239.2	258.0
flooradapt $C(\{\mathbf{H}\})$	138.1	155.0	183.5	191.1	214.2	212.5	222.9

measured [15]), the MSE produced by the rounding operation is approximately $\sigma_h^2 = \frac{1}{12}\Delta_h^2$. This relationship underpins the initialization of the quantization step-sizes. Recall that in this paper the step-sizes are initialized to be less than one, the largest signal value is 39. Successive halving of the interval Δ_h causes the variance to reduce according to: $\sigma_h^2 = \frac{1}{12}\Delta_h^2 \frac{1}{2^s}$. In effect as $\Delta_h \mapsto 0$ the quantization error is reduced and QNMF begins to approximate the original NMF – a related analysis holds for the floor and not-round functions. It is reasonable to assume that in the limit QNMF reduces to original NMF. In short, monotonic convergence is feasible (for a large enough s). Successive halving (of Δ_h) in this manner exhibits favourable complexity. In practice, we have observed that the number of Bernoulli trials required to achieve minimization using the random quantizers is small. Moreover step-size division is rarely required, in fact s is fixed for all our experiments.

VII. CONCLUSIONS

Our first conclusion is that QNMF is more amenable to storage than NMF due to its reduced factor-cardinality. It trades-off only 1-4dB to achieve a reduction of 20-70 (500-1000) times the cardinality of the \mathbf{W} (\mathbf{H}) matrix element set to achieve this reduction. We make the recommendation that based on the cardinality gains achieved by QNMF, efficient and accurate compression-factorization can be achieved using QNMF. A second conclusion is that our empirical convergence evaluation of QNMF substantiates the claim that quantization does not significantly affect convergence. In future work we will give a more formal analysis of QNMF convergence and consider the impact of quantization on domain specific perceptual measures.

REFERENCES

- [1] Ruairí de Fréin, K. Drakakis, and Scott Rickard, “Portfolio diversification using subspace factorizations,” in *Inf. Scien. Sys., 42nd Ann. Conf.*, 2008, pp. 1075–80.
- [2] D. FitzGerald, “User assisted source separation using non-negative matrix factorisation,” *22nd IET Irish Sig. and Sys. Conf.*, 2011, Dublin.

- [3] R. Jaiswal, D. FitzGerald, E. Coyle, and S.T. Rickard, "Shifted NMF using an Efficient Constant-Q Transform for Monaural Sound Source Separation," *22nd IET Irish Sig. and Sys. Conf.*, 2011, Dublin.
- [4] P.D. O'Grady and S.T. Rickard, "Automatic hexaphonic guitar transcription using non-negative constraints," *IET ISSC*, June 2009, Dublin.
- [5] D.D. Lee and S.H. Seung, "Algorithms for Non-negative Matrix Factorization," in *NIPS*. 2000, pp. 556–62, MIT Press.
- [6] D. Donoho and V. Stodden, "When does non-negative matrix factorization give correct decomposition into parts?," in *NIPS*. 2003, MIT Press.
- [7] Ruairí de Fréin, "Ghostbusters: A parts-based NMF algorithm," *Signals and Systems Conference (ISSC 2013), 24th IET Irish*, pp. 1–8, June 2013.
- [8] S. Rickard and A. Cichocki, "When is non-negative matrix decomposition unique?," in *IEEE 42nd CISS*, 2008, pp. 1091–2.
- [9] H. Laurberg, M.G. Christensen, M.D. Plumbley, L.K. Hansen, and S.H. Jensen, "Theorems on Positive Data: On the Uniqueness of NMF," *Comput Intell Neurosci*, 2008.
- [10] G. Tsagkatakis and P. Tsakalides, "Dictionary based reconstruction and classification of randomly sampled sensor network data," in *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2012 IEEE 7th*, June 2012, pp. 117–120.
- [11] L. Kong, M. Xia, X. Liu, G. Chen, Y. Gu, M. Wu, and X. Liu, "Data loss and reconstruction in wireless sensor networks," *Par. Distr. Sys., IEEE Trans.*, , no. 99, 2013.
- [12] T. Andre, M. Antonini, M. Barlaud, and R.M. Gray, "Entropy-based distortion measure and bit allocation for wavelet image compression," *Im. Proc., IEEE Trans.*, vol. 16, no. 12, pp. 3058–3064, Dec 2007.
- [13] D. W E Schobben, R.A. Beuker, and W. Oomen, "Dither and data compression," *Sig. Proc., IEEE Trans.*, vol. 45, no. 8, pp. 2097–2101, Aug 1997.
- [14] V. Goyal, M. Vetterli, and N.T. Thao, "Quantized Overcomplete Expansions in R^N : Analysis, Synthesis and Algorithms," *IEEE Trans. Inf. Th.*, vol. 44, no. 1, Jan 98.
- [15] H. Gish and J. Pierce, "Asymptotically efficient quantizing," *Inf. Th., IEEE Trans.*, vol. 14, no. 5, pp. 676–683, Sep 1968.
- [16] R.M. Gray and D.L. Neuhoff, "Quantization," *IEEE Trans. Inf. Th.*, vol. 44, no. 6, Oct. 1998.
- [17] J Beirlant, E J Dudewicz, L Györfi, and E C Van Der Meulen, "Non-parametric entropy estimation: An overview," *International Journal of Mathematical and Statistical Sciences*, vol. 6, no. 1, pp. 17–39, 1997.