

A Rigorous Link Between Self-Organizing Maps and Gaussian Mixture Models

Alexander Gepperth and Benedikt Pfüll

University of Applied Sciences Fulda, Leipzigerstr.123, 36037 Fulda, Germany
{alexander.gepperth,benedikt.pfuelb}@cs.hs-fulda.de
www.hs-fulda.de

Abstract. This work presents a mathematical treatment of the relation between Self-Organizing Maps (SOMs) and Gaussian Mixture Models (GMMs). We show that energy-based SOM models can be interpreted as performing gradient descent, minimizing an approximation to the GMM log-likelihood that is particularly valid for high data dimensionalities. The SOM-like decrease of the neighborhood radius can be understood as an annealing procedure ensuring that gradient descent does not get stuck in undesirable local minima. This link allows to treat SOMs as generative probabilistic models, giving a formal justification for using SOMs, e.g., to detect outliers, or for sampling.

Keywords: Self-Organizing Maps · Gaussian Mixture Models · Stochastic Gradient Descent

1 Introduction

This theoretical work is set in the context of unsupervised clustering and density estimation methods and establishes a mathematical link between two important representatives: Self-Organizing Maps (SOMs, [1, 10]) and Gaussian Mixture Models (GMMs, [2]), both of which have a long history in machine learning. There are significant overlaps between SOMs and GMMs, and both models have been used for data visualization and outlier detection. They are both based on Euclidean distances and model data distributions by prototypes or centroids. At the same time, there are some differences: GMMs, as fully generative models with a clear probabilistic interpretation, can additionally be used for sampling purposes. Typically, GMMs are trained batch-wise, repeatedly processing all available data in successive iterations of the Expectation-Maximization (EM) algorithm. In contrast, SOMs are trained online, processing one sample at a time. The training of GMMs is based on a loss function, usually referred to as *incomplete-data log-likelihood* or just log-likelihood. Training by Stochastic Gradient Descent (SGD) is possible, as well, although few authors have explored this [9]. SOMs are not based on a loss function, but there are model extensions [4, 7] that propose a simple loss function at the expense of very slight differences in model equations. Lastly, GMMs have a simple probabilistic interpretation

as they attempt to model the density of observed data points. For this reason, GMMs may be used for outlier detection, clustering and, most importantly, sampling. In contrast to that, SOMs are typically restricted to clustering and visualization due to the topological organization of prototypes, which does not apply to GMMs.

1.1 Problem Statement

SOMs are simple to use, implement and visualize, and, despite the absence of theoretical guarantees, have a very robust training convergence. However, their interpretation remains unclear. This particularly concerns the probabilistic meaning of input-prototype distances. Different authors propose using the Best-Matching Unit (BMU) position only, while others make use of the associated input-prototype distance, or even the combination of all distances [6]. Having a clear interpretation of these quantities would help researchers tremendously when interpreting trained SOMs. The question whether SOMs actually perform density estimation is important for justifying outlier detection or clustering applications. Last but not least, a probabilistic interpretation of SOMs, preferably a simple one in terms of the well-known GMMs, would help researchers to understand how sampling from SOMs can be performed.

1.2 Results and Contribution

This article aims at explaining SOM training as Stochastic Gradient Descent (SGD) using an energy function that is a particular approximation to the GMM log-likelihood. SOM training is shown to be an approximation to training GMMs with tied, spherical covariance matrices where constant factors have been discarded from the probability computations. This identification allows to interpret SOMs in a probabilistic manner, particularly for:

- outlier detection (not only the position of the BMU can be taken into account, but also the associated input-prototype distance since it has a probabilistic interpretation) and
- sampling (understanding what SOM prototypes actually represent, it is possible to generate new samples from SOMs with the knowledge that this is actually sanctioned by theory)

1.3 Related Work

Several authors have attempted to establish a link between SOMs and GMMs. In [8], an EM algorithm for SOMs is given, emphasizing the close links between both models. Verbeek et al. [14] emphasizes that GMMs are regularized to show a SOM-like behavior of self-organization. A similar idea of component averaging to obtain SOM-like normal ordering and thus improved convergence was previously demonstrated in [11]. An energy-based extension of SOMs suggesting a close relationship to GMMs is given in [7], with an improved version described in [4]. So far, no scientific work has tried to explain SOMs as an approximation to GMMs in a way that is comparable to this work.

2 Main Proof

The general outline of proof is depicted in Fig. 1. We start with a description of GMMs in Sec. 2.1. Subsequently, a transition from exact GMMs to the popular Max-Component (MC) approximation of its loss is described in Sec. 2.2. Then, we propose a SOM-inspired annealing procedure for performing optimization of approximate GMMs in Sec. 2.3 and explain its function in the context of SGD. Finally, we show that this annealed training procedure is equivalent to training energy-based SOMs in Sec. 2.4, which are a faithful approximation of the original SOM model, outlined in Sec. 2.5.




Fig. 1: Outline of main proof.

2.1 Default GMM Model, Notation and GMM Training

Gaussian Mixture Models (GMMs) are probabilistic latent-variable models that aim to explain the distribution of observed data samples $X = \{\mathbf{x}_n\}$. It is assumed that samples generated from a known parametric distribution, which depends on parameters θ and unobservable latent variables $Z = \{\mathbf{z}_n \in \mathbb{R}^K\}$, $z_{nk} \in \{0, 1\}$ and $\sum_k z_{nk} = 1$. The *complete-data probability* reads

$$p(X, Z) = \prod_n \prod_k [p_k(\mathbf{x}_n)]^{\pi_k z_{nk}} \quad (1)$$

where mixture components $p_k(\mathbf{x}_n)$ are modeled as multi-variate Gaussians, and whose parameters are the centroids $\boldsymbol{\mu}_k$ and the positive-definite covariance matrices Σ_k (both omitted from the notation for conciseness). $p_k(\mathbf{x}_n)$ represents the probability of observing the data vector \mathbf{x}_n , if sampled from mixture component k . The number of Gaussian mixture components K is a free parameter of the model, and the component weights π_k must have a sum of 1. Since the latent variables are not observable, it makes sense to marginalize them out by summing over the discrete set of all possible values $\mathcal{Z} = \{\mathbf{e}_j | j = 1 \dots K\}$, giving

$$p(X) = \prod_n \sum_{\mathbf{z}_n \in \mathcal{Z}} \prod_k p_k(\mathbf{x}_n)^{\pi_k z_{nk}}. \quad (2)$$

Taking the logarithm and normalizing by the number of samples (N) provides the *incomplete-data log-likelihood*, which contains only observable quantities and is, therefore, a suitable starting point for optimization:

$$\mathcal{L} = \frac{1}{N} \sum_{n=1}^N \log \sum_k \pi_k p_k(\mathbf{x}_n). \quad (3)$$

Concise Problem Statement

When training GMMs, one aims at finding parameters $\boldsymbol{\mu}_k, \Sigma_k$ that (locally) maximize Eq. (3). This is usually performed by using a procedure called Expectation-Maximization (EM, [2, 5]) which is applicable to many latent-data (mixture) models. Of course, a principled alternative to EM is an approach purely based on batches or SGD, the latter being an approximation justified by the Robbins-Monro procedure [12]. In this article, we will investigate how SGD optimization of Eq. (3) can be related to the training of SOMs.

Respecting GMM Constraints in SGD

GMMs impose the following constraints on the parameters $\pi_k, \boldsymbol{\mu}_k$ and Σ_k :

- weights must be normalized: $\sum_k \pi_k = 1$
- covariance matrices must be positive-definite: $\mathbf{x}^T \Sigma_k \mathbf{x} \geq 0 \forall k, \mathbf{x}$

The first constraint can be enforced after each gradient decent step by setting $\pi_k \rightarrow \frac{\pi_k}{\sum_j \pi_j}$. For the second constraint, we consider diagonal covariance matrices only, which is sufficient for establishing a link to SOMs. A simple strategy in this setting is to re-parameterize covariance matrices Σ_k by their inverse (denoted as precision matrices) $\mathbf{P}_k = \Sigma^{-1}$. We then re-write this as $\mathbf{P}_k = \mathbf{D}_k \mathbf{D}_k$, which ensures positive-definiteness of \mathbf{P}, Σ . The diagonal entries s_k of Σ_k can thus be re-written as $s_{ki} = d_{ki}^{-2}$, whereas \mathbf{d}_k are the diagonal entries of \mathbf{D}_k .

2.2 Max-Component Approximation

In Eq. (3), we observe that the component weights π_k and the conditional probabilities $p(\mathbf{x})$ are positive by definition. It is, therefore, evident that any single component of the inner sum over the components k is a lower bound of the entire inner sum. The largest of these K lower bounds is given by the maximum of the components, so that it results in

$$\begin{aligned} \mathcal{L} &= \frac{1}{N} \sum_{n=1}^N \log \sum_k \pi_k p(\mathbf{x}_n) \leq \hat{\mathcal{L}} = \frac{1}{N} \sum_n \log \max_k (\pi_k p(\mathbf{x}_n)) \\ &= \frac{1}{N} \sum_n \max_k \log (\pi_k p(\mathbf{x}_n)). \end{aligned} \tag{4}$$

Eq. (4) displays what we refer to as *Max-Component approximation* to the log-likelihood. Since $\hat{\mathcal{L}} \leq \mathcal{L}$, we can increase \mathcal{L} by maximizing $\hat{\mathcal{L}}$. The advantage of $\hat{\mathcal{L}}$ is that it is not affected by numerical instabilities the way \mathcal{L} is. Moreover, it breaks the symmetry between mixture components, thus, avoiding degenerate local optima during early training. Apart from facilitating the relation to SOMs, this is an interesting idea in its own right, which was first proposed in [3].

Undesirable Local Optima

GMMs are usually trained using EM after a k-means initialization of the centroids. Since this work explores the relation to SOMs, which are mainly trained from scratch, we investigate SGD-based training of GMMs without k-means initialization. A major problem in this setting are undesirable local optima, both for the full log-likelihood \mathcal{L} and its approximation $\hat{\mathcal{L}}$. To show this, we parameterize the component probabilities by the precision matrices $\mathbf{P}_k = \Sigma_k^{-1}$ and compute

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \boldsymbol{\mu}_k} &= \mathbb{E}_n [\mathbf{P}_k (\mathbf{x}_n - \boldsymbol{\mu}_k) \gamma_{nk}] \\ \frac{\partial \mathcal{L}}{\partial \mathbf{P}_k} &= \mathbb{E}_n [((\mathbf{P}_k)^{-1} - (\mathbf{x}_n - \boldsymbol{\mu}_k)(\mathbf{x}_n - \boldsymbol{\mu}_k)^T) \gamma_{nk}] \\ \frac{\partial \mathcal{L}}{\partial \pi_k} &= \pi_k^{-1} \mathbb{E}_n [\gamma_{nk}] \end{aligned} \quad (5)$$

whereas $\gamma_{nk} \in [0, 1]$ denote standard GMM responsibilities given by

$$\gamma_{nk} = \frac{\pi_k p_k(\mathbf{x}_n)}{\sum_k \pi_k p_k(\mathbf{x}_n)}. \quad (6)$$

Degenerate Solution This solution universally occurs when optimizing \mathcal{L} by SGD, and represents an obstacle for naive SGD. All components have the same weight, centroid and covariance matrix: $\pi_k \approx \frac{1}{K}$, $\boldsymbol{\mu}_k = \mathbb{E}[\mathbf{X}]$, $\Sigma_k = \text{Cov}(\mathbf{X}) \forall k$. Since the responsibilities are now uniformly $1/K$, it results from Eq. (5) that all gradients vanish. This effect is avoided by $\hat{\mathcal{L}}$ as only a subset of components is updated by SGD, which breaks the symmetry of the degenerate solution.

Single/Sparse-Component Solution Optimizing $\hat{\mathcal{L}}$ by SGD, however, leads to another class of unwanted local optima: A single component k^* has a weight close to 1, with its centroid and covariance matrix being given by the mean and covariance of the data: $\pi_{k^*} \approx 1$, $\boldsymbol{\mu}_{k^*} = \mathbb{E}[\mathbf{X}]$, $\Sigma_{k^*} = \text{Cov}(\mathbf{X})$. For $\hat{\mathcal{L}}$, the gradients in Eq. (5) stay the same except for $\gamma_{nk} = \delta_{nk^*}$ from which we conclude that the gradient w.r.t. \mathbf{P}_k and $\boldsymbol{\mu}_k$ vanishes $\forall k$. The gradient w.r.t. π_k does not vanish, but is δ_{kk^*} , which disappears after enforcing the normalization constraint (see Sec. 2.1). A variant is the sparse-component solution where only a few components have non-zero weights, so that the gradients vanish for the same reasons.

2.3 Annealing Procedure

A simple SOM-inspired approach to avoid these undesirable solutions is to punish their characteristic response patterns by an appropriate modification of the (approximate) loss function that is maximized, i.e., $\hat{\mathcal{L}}$. We introduce what we call *smoothed Max-Component log-likelihood* $\hat{\mathcal{L}}^\sigma$, inspired by SOM training:

$$\hat{\mathcal{L}}^\sigma = \frac{1}{N} \sum_n \max_k \left(\sum_j \mathbf{g}_{kj} \log(\pi_j p(\mathbf{x}_n)) \right). \quad (7)$$

Here, we assign a normalized coefficient vector \mathbf{g}_k to each Gaussian mixture component k . The entries of \mathbf{g}_k are computed in the following way:

- Assume that the K Gaussian components are arranged on a 1D grid of dimensions $(1, K)$ or on a 2D grid of dimensions (\sqrt{K}, \sqrt{K}) . As a result, each linear component index k has a unique associated 1D or 2D coordinate $\mathbf{c}(k)$.
- Assume that the vector \mathbf{g}_k of length K is actually representing a 1D structure of dimension $(1, K)$ or a 2D structure of dimension (\sqrt{K}, \sqrt{K}) . Each linear vector index j in \mathbf{g}_k has a unique associated 1D or 2D coordinate $\mathbf{c}(j)$.
- The entries of the vector \mathbf{g}_k are computed as

$$g_{kj} = \exp\left(-\frac{(\mathbf{c}(j) - \mathbf{c}(k))^2}{2\sigma^2}\right) \quad (8)$$

and subsequently normalized to have a unit sum. Essentially, Eq. (7) represents a convolution of $\log \pi_k p_k(\mathbf{x})$, arranged on a periodic 2D grid with a Gaussian convolution filter, resulting in a smoothing operation. The 2D variance σ in Eq. (8) is a parameter that must be set as a function of the grid size so that Gaussians are neither homogeneous, nor delta peaks. Hence, the loss function in Eq. (7) is maximized if the log probabilities follow an uni-modal Gaussian profile of variance σ , whereas single-component solutions are punished.

It is trivial to see that the annealed loss function in Eq. (7) reverts to the non-annealed form Eq. (4) in the limit where $\sigma \rightarrow 0$. This is due to the fact that vectors \mathbf{g}_k approach Kronecker deltas in this case with only a single entry of value 1. Thereby, the inner sum in Eq. (7) is removed. By making σ time-dependent in a SOM-like manner, starting at a value of $\sigma(t_0) \equiv \sigma_0$ and then reducing it to a small final value $\sigma(t_\infty) \equiv \sigma_\infty$. The result is a smooth transition from the annealed loss function Eq. (7) to the original max-component log-likelihood Eq. (4). Time dependency of $\sigma(t)$ can be chosen to be:

$$\sigma(t) = \begin{cases} \sigma_0 & t < t_0 \\ \sigma_\infty & t > t_\infty \\ \sigma_0 \exp(-\tau t) & t_0 < t < t_\infty \end{cases} \quad (9)$$

where the time constant in the exponential is chosen as $\tau = \log \frac{\sigma_0 - \sigma_\infty}{t_\infty - t_0}$ to ensure a smooth transition. This is quite common while training SOMs where the neighborhood radius is similarly decreased.

2.4 Link to Energy-based SOM Models

The standard Self-Organizing Map (SOM) has no energy function that is minimized. However, some modifications (see [4, 7]) have been proposed to ensure the existence of a C^∞ energy function. These energy-based SOM models reproduce all features of the original model and use a learning rule that the original SOM algorithm is actually approximating very closely. In the notation of this article,

SOMs model the data through K prototypes $\boldsymbol{\mu}_k$ and K neighborhood functions g_k defined on a periodic 2D grid. Their energy function is written as

$$\mathcal{L}_{SOM} = \frac{1}{N} \sum_n \min_k \sum_j g_{kj} \|\mathbf{x}_n - \boldsymbol{\mu}_j\|^2, \quad (10)$$

whose optimization by SGD initiates the learning rule for energy-based SOMs:

$$\boldsymbol{\mu}_k(t+1) = \boldsymbol{\mu}_k + \epsilon g_{ki^*}(\mathbf{x} - \boldsymbol{\mu}_k) \quad (11)$$

with the Best-Matching Unit (BMU) having index i^* . In contrast to the standard SOM model the BMU is determined as $i^* = \operatorname{argmax}_i \sum_j g_{ij} \|\mathbf{x} - \boldsymbol{\mu}_i\|$. The link to the standard SOM model is the observation that for small values of the neighborhood radius $\sigma(t)$ the convolution vanishes and the original SOM learning rule is recovered. This is typically the case after the model has initially converged (sometimes referred to as “normal ordering”).

2.5 Equivalence to SOMs

When writing out $\log \pi_k p_k(\mathbf{x}) = -\sum_j \frac{d_{kj}^2}{2} (\mathbf{x}_j - \boldsymbol{\mu}_{kj})$, tying the variances so that $d_{kj} = d \forall j$ and fixing the weights to $\pi_k = \frac{1}{K}$ in Eq. (7) we find that the energy function Eq. (10) becomes

$$\hat{\mathcal{L}}^\sigma = \frac{1}{N} \sum_n \max_k \left(\sum_j g_{kj} \left(-\log K - \frac{d^2}{2} \|\mathbf{x}_n - \boldsymbol{\mu}_j\| \right) \right) \quad (12)$$

$$= \frac{d^2}{2N} \sum_n \min_k \sum_j g_{kj} \|\mathbf{x}_n - \boldsymbol{\mu}_j\| + \text{const.} \quad (13)$$

In fact Eq. (7) is identical to Eq. (10), except for a constant factor that can be discarded and a scaling factor defined by the common tied precision d^2 . The minus sign just converts the max into a min operation, as distances and precisions are positive. The annealing procedure of Eq. (9) is identical to the method for reducing the neighborhood radius during SOM training as well.

Energy-based SOMs are a particular formulation (tied weights, constant spherical variances) of GMMs which are approximated by a commonly accepted method. Training energy-based SOMs in the traditional way results in the optimization of GMMs by SGD, where training procedures are, again, identical.

3 Experiments

This section presents a simple proof-of-concept that the described SGD-based training scheme is indeed practical. It is not meant to be an exhaustive empirical proof and is, therefore, just conducted on a common image dataset. For this experiment, we use the MADBase dataset [13] containing 60 000 grayscale images

of handwritten arabic digits in a resolution of 28×28 pixels. The number of training iterations $T = 24\,000$ with a batch size of 1, as it is common for SOMs. We use a GMM with $K = 25$ components, whose centroids are initialized to small random values, whereas weights are initially equiprobable and precisions are uniformly set to $d^2 = 5$ (we found that precisions should initially be as large as possible). We set $t_0 = 0.3T$ and stop at $t_\infty = 0.8T$ (proportional to the maximum number of iterations). σ_0 starts out at 1.2 (proportional to the map size) and is reduced to a value of 0.01. The learning rate is similarly decayed to speed up convergence, although this is not a requirement, with $\epsilon_0 = 0.05$ and $\epsilon_\infty = 0.009$. In Fig. 2, three states of the trained GMM are depicted (at iteration $t_1 = 3120$, $t_2 = 8640$ and $t_3 = 22080$).




Ensemble of “neighborhood functions” g_{kj} Centroids μ_k Trend of $\sigma(t)$ and $\epsilon(t)$

Fig. 2: Visualization of GMM evolution at three points in time during training (from left to right: iterations $t_1 = 3120$, $t_2 = 8640$ and $t_3 = 22080$).

We observe from the development of the prototypes that GMM training converges, and that the learned centroids are representing the dataset well. It can also be seen that prototypes are initially blurred and get refined over time, which resembles SOMs. No undesired local optima were encountered when this experiment was repeated 100 times.

4 Discussion

Approximations The approximations on the way from GMMs to SOMs are Eq. (4) and the approximations made by energy-based SOM models (see [8] for a discussion). It is shown in [3] that the quality of the first approximation is excellent, since inter-cluster distances tend to be large, and it is more probable that a single component can explain the data.

Consequences The identification of SOMs as special approximation of GMMs allows for the performance of typical GMM functions (sampling, outlier detection) with trained SOMs. The basic quantity to be considered here is the input-prototype distance of the Best-Matching Unit (BMU) since it corresponds to a log probability. In particular, the following consequences were to be expected:

- **Outlier Detection** The value of the smallest input-prototype distance is the relevant one, as it represents $\hat{\mathcal{L}}$ for a single sample, which in turn approximates the incomplete-data log-likelihood \mathcal{L} . In practice, it can be advantageous to average $\hat{\mathcal{L}}$ over several samples to be robust against noise.
- **Clustering** SOM prototypes should be viewed as cluster centers, and inputs can be assigned to the prototype with the smallest input-prototype distance.
- **Sampling** Sampling from SOMs should be performed in the same way as from GMMs. Thus, in order to create a sample, a random prototype has to be selected first (since weights are tied no multinomials are needed here). Second, a sample from a Gaussian distribution with precision 2 and centered on the selected prototype along all axes needs to be drawn.

4.1 Summary and Conclusion

To our knowledge, this is the first time that a rigorous link between SOMs and GMMs has been established, based on a comparison of loss/energy functions. It is thus shown that SOMs actually implement an annealing-type approximation to the full GMM model with fixed component weights and tied diagonal variances. To give more weight to the mathematical proof, we validate the SGD-based approach to optimize GMMs in practice.

References

1. Cottrell, M., Fort, J.C., Pagès, G.: Two or three things that we know about the Kohonen algorithm. *ESANN '1999 proceedings* pp. 235–244 (1994)
2. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum Likelihood from Incomplete Data Via the EM Algorithm , vol. 39. Wiley Online Library (1977). <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>
3. Dognin, P.L., Goel, V., Hershey, J.R., Olsen, P.A.: A fast, accurate approximation to log likelihood of Gaussian mixture models. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings* **3**, 3817–3820 (2009). <https://doi.org/10.1109/ICASSP.2009.4960459>
4. Gepperth, A.: An energy-based SOM model not requiring periodic boundary conditions. *Neural Computing and Applications* (2019). <https://doi.org/10.1007/s00521-019-04028-9>
5. Hartley, H.O.: Maximum likelihood estimation from incomplete data. In: *Biometrics*. vol. 14, pp. 174–194 (1958)
6. Hecht, T., Lefort, M., Gepperth, A.: Using self-organizing maps for regression: The importance of the output function. *23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN - Proceedings* pp. 107–112 (2015)
7. Heskes, T.: Energy functions for self-organizing maps. *Kohonen Maps* pp. 303–315 (1999). <https://doi.org/10.1016/b978-044450270-4/50024-3>
8. Heskes, T.: Self-organizing maps, vector quantization, and mixture modeling. *IEEE Transactions on Neural Networks* **12**(6), 1299–1305 (2001). <https://doi.org/10.1109/72.963766>
9. Hosseini, R., Sra, S.: Matrix manifold optimization for Gaussian mixtures. *Advances in Neural Information Processing Systems* pp. 910–918 (2015)
10. Kohonen, T.: The Self-organizing Map. *Proceedings of the IEEE* **78**(9), 1464–1480 (1990)
11. Ormoneit, D., Tresp, V.: Averaging, maximum penalized likelihood and Bayesian estimation for improving Gaussian mixture probability density estimates. *IEEE Transactions on Neural Networks* **9**(4), 639–650 (1998). <https://doi.org/10.1109/72.701177>
12. Robbins, H., Monro, S.: A Stochastic Approximation Method. *The Annals of Mathematical Statistics* **22**(3), 400–407 (1951). <https://doi.org/10.1214/aoms/1177729586>
13. Sherif, A., Ezzat, E.S.: Ahdbase, <http://datacenter.aucegypt.edu/shazeem/>
14. Verbeek, J.J., Vlassis, N., Kröse, B.J.: Self-Organizing Mixture Models. *Neurocomputing* **63**, 99–123 (2005). <https://doi.org/10.1016/j.neucom.2004.04.008>