

An Information-Theoretic Framework for Semantic-Multimedia Indexing

João Magalhães
Department of Computing
Imperial College London
London, UK

j.magalhaes@imperial.ac.uk

Stefan R uger
Knowledge Media Institute
The Open University
Milton Keynes, UK

s.rueger@open.ac.uk

Abstract.

To solve the problem of indexing collections with diverse text documents, image documents, or documents with both text and images, one needs to develop a model that supports heterogeneous types of documents. In this paper, we show how information theory supplies us with the tools necessary to develop a unique model for text, image, and text/image retrieval. In our approach, for each possible query keyword we estimate a maximum entropy model based on exclusively continuous features that were pre-processed. The unique continuous feature-space of text and visual data is constructed by using a minimum description length criterion to find the optimal feature-space representation (optimal from an information theory point of view). We evaluate our approach in three experiments: only text retrieval, only image retrieval, and text combined with image retrieval.

1 Introduction

Demand for techniques that handle both text and image based documents is increasing with the wide spread of search applications. It is impossible to conceive nowadays a world without systems that allow us to search for specific news articles, scientific papers, or information in general. Users want more: they want to submit the same query to search for text documents, visual documents, or documents with both media, e.g., photographs with captions, video shots (key-frames and speech). To achieve this, a new breed of information retrieval models is required: one that seamlessly integrates heterogeneous data. Thus, in this thesis we assume that in any given collection \mathcal{D} of N multimedia documents

$$\mathcal{D} = \{d^1, d^2, \dots, d^N\}, \quad (1)$$

each document is characterized by a vector

$$d^j = (d_T^j, d_V^j, d_W^j), \quad (2)$$

composed by a feature vector d_T describing the text part of the document, a feature vector d_V describing the visual part of the document, and a keyword vector d_W describing the semantics of the document. Text based features such as bag-of-words, part-of-speech and named entities, and low-level

visual features such as texture, colour and shape are widely studied, and several algorithms exist to extract these types of low-level feature vectors. Keyword features representing multimedia information have a less consensual solution because of the ambiguity and subjectivity of the information that they try to describe – the semantic content of a multimedia document. The semantic description of multimedia information, the feature vector d_W , is the core topic of this thesis. To describe the semantics of multimedia information we define the set

$$\mathcal{W} = \{w_1, \dots, w_L\} \quad (3)$$

as a vocabulary of L keywords. These keywords are linguistic representations of abstract or concrete concepts that we want to detect in multimedia documents. The feature vector d_W is formally defined as

$$d_W^j = (d_{W,1}^j, d_{W,2}^j, \dots, d_{W,L}^j) \quad (4)$$

where each component $d_{W,t}^j$ is a score indicating the confidence that keyword w_t is present in that particular document. The concepts may not be explicitly present in multimedia information, methods are required to compute the likelihood that the keyword is actually present in the multimedia document.

Equation (2) shows us the other information that we have about documents: text and visual feature. Thus, to compute the components of the keyword vector d_W^j we shall use text and visual feature data. This leads us to the definition of each component of the keyword vector as

$$d_{W,t}^j = p(y_t^j = 1 \mid d_T^j, d_V^j), \quad (5)$$

where the random variable $y_t^j = \{1, 0\}$ indicates the presence/not-presence of keyword w_t on document d^j given its text feature vector d_T^j and visual feature vector d_V^j . This enables the semantic indexing of multimedia content which allows users to *submit the same query to search for text documents, visual documents, or documents with both media, e.g., photographs with captions, video shots (key-frames and speech)*. Equation (2) integrates heterogeneous representations of a multimedia document (text, image and semantic) and Equation (5) will make multimedia information searchable with the same type of queries for all type of media.

In this paper we address the problem of estimating a statistical model for Equation (5). We shall

propose a statistical framework that can simultaneously model text-only documents, image-only documents, and documents with both text and images.

2 An Information-Theoretic Framework

Our first objective is to compute the components of keyword feature vectors d_W representing the semantics of multimedia documents. For this we will estimate and select a model $\beta_t \in \Theta$, from a set Θ of candidate models, that best represents the keyword w_t in terms of text data and visual data. We omit model β_t of keyword w_t from Equation (5) for notational simplicity. The expression can now be written as:

$$d_{W,t}^j = p\left(y_t^j = 1 \mid d_T^j, d_V^j, \beta_t\right) \quad (6)$$

The statistical model $\beta_t \in \Theta$ can assume many forms (e.g., nearest neighbour, neural networks, linear models, support vector machines) according to the family of algorithms and to the complexity of the specific algorithm within a particular family of algorithms. The choice of the family of algorithms is done by examining the requirements that multimedia information retrieval applications face in a real world scenario:

- Arbitrary addition and removal of keywords
- Easy update of existing keyword models with new training data
- Seamless integration of heterogeneous types of data
- Computationally efficient indexing of multimedia information
- Good retrieval effectiveness

The first two requirements concern an important practical aspect in large-scale multimedia indexes – the integrity of the index when keyword models are modified. When a keyword model is modified (added, removed or updated) the index can be affected in two ways: if keyword models are dependent then the entire index becomes obsolete; if keyword models are independent then only the part of the index concerning that keyword becomes obsolete. This leads to a solution where keyword models are independent so that a modification in one keyword model will have a minor influence on the indexes. Thus, presence of keywords shall be represented by Bernoulli random variables

$$y_t \sim \text{Bernoulli}(p_t), \quad (7)$$

where p_t is the probability of keyword w_t .

The remaining three requirements can be difficult to accommodate in a unique model: support of multi-modal information, be able to quickly index new multimedia content and to achieve a good accuracy. When modelling multi-modal keywords, one has to deal with both dense feature spaces and sparse features spaces. On one hand visual feature data can be very dense making its modelling difficult due to the irregular frontiers caused by concept cross-interference. Expanding the original feature space into higher-dimensional ones results in a sparser feature space where the modelling of the data can be made easier. On the other hand, text feature spaces are typically too sparse making its modelling difficult because there is not enough support data to estimate the details of concept models. In these situations we have to compress the feature space into a lower dimensional space where data is compressed into a more dense space. These transformations of the original feature space into a space where the data is optimally distributed is represented as

$$F(d_T^j, d_V^j) = (F_T(d_T^j), F_V(d_V^j)), \quad (8)$$

where $F_T(d_T^j)$ correspond to the text data transformation and $F_V(d_V^j)$ correspond to the visual data transformation. This renders the final expression for the components of keyword feature vectors as

$$d_{W,t}^j = p(y_t^j = 1 | F(d_T^j, d_V^j), \beta_t). \quad (9)$$

The transformation of multimedia document features only need to be computed once for all keyword models, in other words, the transformation is independent of the keyword models. The interesting

implication of this fact is that it can reduce the indexing computational complexity: because the transformation generates a high-dimensional space, one can limit the keyword model search space Θ to the family of linear models which have a very low computational complexity in the classification phase (but not necessarily in the learning phase). Besides the low computational complexity, linear models offer other interesting advantages: support of high-dimensional data (easy integration of heterogeneous data), naturally embedded background knowledge in the form of priors (ideal for keyword model update) and good accuracy (retrieval effectiveness).

In the remainder of Chapter 4 I will present and propose $F_T(d_T)$ and $F_V(d_V)$ the transformations of visual and text data. Chapter 5 presents several linear models β_i to represent keywords.

3 Related Work

In text retrieval the search process is triggered by a text query that can be compared directly to the corpus of the documents in the collection. Since we want to offer a common query interface for both text and images we need to define a common vocabulary of keywords to query all possible types of documents. Therefore the present work is related to text categorization, image annotation and multimodal content annotation. We will now look at these three areas with a view to seamlessly integrate text and image data on the same framework.

Text categorization models pre-process data by removing stop-words and rare words, stemming, and finally term-weighting. Due to the high-dimensional feature space of text data most text categorization algorithms are linear models such as naïve Bayes [28], maximum entropy [30], Support Vector Machines [19], regularized linear models [46], and Linear Least Squares Fit [42]. Joachims [19] applies SVMs directly to the text terms. Text is ideal for applying SVMs without the need of a kernel function because data is already sparse and high-dimensional. Linear models fitted by least squares such as the one by Yang et. al [42] offer good precision, and in particular regularized linear methods, such as the one we propose and the one by Zhang and Oles [46], perform similarly to SVMs, with the advantage of yielding a probability density model. The maximum entropy classification model proposed by Nigam [30] defines a set of features that are dependent on the class being evaluated while we use a unique set of features for all keywords. The proposed maximum entropy framework has the same characteristics and performance of

linear models (logistic regression, least squares) with the crucial advantage that while these approaches have no automatic mechanism to select a vocabulary size we use the minimum description length principle to select its optimal size.

Yang [41], and Yang and Liu [43] have compared a number of text classification algorithms and reported their performances on different text collections. Their results indicate that k-Nearest Neighbour, SVMs, and LLSF are the best classifiers. Note that nearest neighbour approaches have certain characteristics (see [14]) that make them computationally too complex to handle large-scale indexing.

The simplest image annotation models deploy a traditional multi-class supervised learning model and learn the class-conditional probability density distribution of each keyword w given its training data x . Bayes law is used to model $p(x | w)$, the features data density distribution of a given keyword. Several techniques to model $p(x | w)$ with different types of probability density distributions have been proposed: Yavlinsky et al. [45] deployed a nonparametric distribution; Carneiro and Vasconcelos [7] a semi-parametric density estimation; Westerveld and de Vries [39] a finite-mixture of Gaussians; while Vailaya et al. [38] apply a vector quantization technique. Density based approaches are among the most successful ones. However, density distributions are not adequate for text because the density models do not get enough support from such sparse data.

Other types of approaches are based on a translation model between keywords and images (global, tiles or regions). Inspired by automatic text translation research, Duygulu et al. [10] developed a method of annotating images with words. First, regions are created using a segmentation algorithm like normalised cuts. For each region, features are computed and then blobs are generated by clustering the image features for these regions across an image collection. The problem is then formulated as learning the correspondence between the discrete vocabulary of blobs and the image keywords. Following the same translation approach [11, 16, 21] have developed a series of translation models that use different models for keywords (multinomial/binomial) and images representations (hard clustered regions, soft clustered regions, tiles).

Hierarchical models have also been used in image annotation such as Barnard and Forsyth's [3] generative hierarchical aspect model inspired by a hierarchical clustering/aspect model. The data are assumed to be generated by a fixed hierarchy of nodes with the leaves of the hierarchy corresponding to

soft clusters. Blei and Jordan [6] propose the correspondence latent Dirichlet allocation model; a Bayesian model for capturing the relations between regions, words and latent variables. The exploitation of hierarchical structures (either of the data or of the parameters) increases the number of parameters (model complexity) to be estimated with the same amount of training data.

Maximum entropy models have also been applied to image annotation [2, 17] and object recognition [22]. All these three approaches have specific features for each class (keywords in our case) which increases the complexity of the system. It is curious to note the large difference in precision results between [17] and [2], we believe that it is related to the lack of regularization and to a differing number of features. These approaches were not as successful as density estimation based models as maximum entropy works best in a high-dimensional sparse feature spaces. The proposed maximum entropy framework tackles this problem by expanding the feature space in a similar spirit to Hoffman's probabilistic Latent Semantic Indexing [15].

These single-modality based approaches are far from our initial goal but by analysing them we can see which family of models can be used to simultaneously model text, image, and multi-modal content. Each modality captures different aspects of that same reality, thus carrying valuable information about each keyword of the vocabulary. The simplest approach to multi-modal analysis is to design a classifier per modality and combine the output of these classifiers. Westerveld, et al. [40] combine the visual model and the text model under the assumption that they are independent, thus the probabilities are simply multiplied. Naphade and Huang [29] model visual features with Gaussian Mixtures Models (GMM), audio features with Hidden Markov Models (HMM) and combine them in a Bayesian network.

In multimedia documents the different modalities contain co-occurring patterns that are synchronised/related in a given way because they represent the same reality. Synchronization/relation and the strategy to combine the multi-modal patterns is a key point of the Semantic pathfinder system proposed by Snoek et al. [36, 37]. Their system uses a unique feature vector that concatenates a rich set of visual features, text features from different sources (ASR, OCR), and audio features. Three types of classifiers are available: logistic regression (which without regularization is known to over-fit [8]), Fisher linear discriminant, and SVMs (offering the best accuracy). The fusion of the different modalities is possible to be done at different levels and it is chosen by cross-validation for each concept. The extremely

high computational complexity required to compute the visual features and to iteratively select the best classifier, the best type of fusion, and the SVMs parameter optimization are serious drawbacks of this system. IBM’s Marvel system [1] has a similar architecture with different learning algorithms to analyse the semantics of multimedia content. These two approaches offer the best performance on the TRECVID2005 conference. Both approaches combine the high-dimensional sparse text features and the low-dimensional dense features on the same feature vector. This might represent a problem for the optimization procedure because the information present on each dimension can be very different. Ideally each dimension should contain the same amount of information and the data density/sparseness should be similar across the entire feature space. The first step of our framework aims at finding this optimal trade-off point by compressing the text feature space dimension and by expanding the visual feature space dimension.

4 Optimal Data Representation

The transformations $F_T(d_T)$ and $F_V(d_V)$ change the representation of the original text and visual feature spaces. As mentioned, transformations $F_T(d_T)$ and $F_V(d_V)$ will adopt specific strategies adequate to the characteristics of each type of data. However, in both cases there is the problem of selecting the optimal transformation from the large number of possible transformations and their varying complexities. In practice, the selection of the optimal transformation is equivalent to old questions like “*how many text features?*” and “*how many visual clusters?*” that are usually addressed by some heuristic method. In this section I shall formally address this problem.

The proposed feature space transformations are inspired by information theory: the space transformation F can be seen as a codebook composed by a set of $M = M_T + M_V$ codewords representing the data space. Given the codebook of a feature space one is able to represent all samples of that feature space as a linear combination of keywords from that codebook. Information theory [9] provides us with a set of information measures that not only assess the amount of information that one single source of data contains, but also the amount of information that two (or more) sources of data have in common. Thus, we employ the minimum description length criterion [33], to infer the optimal

complexity M_T and M_V of each feature space transformation $F_T(d_T)$ and $F_V(d_V)$. Note that I use the word “*optimal*” from an information theory point of view. The treatment of the model selection problem presented in this section is based on [14] and [25].

4.1 Assessing the Data Representation Error

The process of changing the original feature-space representation into the new representation with a given candidate transformation \hat{F} has an associated error. If we represent \hat{F} as the estimated transformation, and G as the lossless transformation that we are trying to estimate, we can compute the mean-squared deviation between the estimated model and the desired response as the error

$$\begin{aligned} \text{Err}_{\mathcal{D}}(d) &= \mathbb{E}\left[\left(G(d) - \hat{F}(d)\right)^2\right] \\ &= \sigma_e^2 + \left(\mathbb{E}[\hat{F}(d)] - G(d)\right)^2 + \mathbb{E}\left[\hat{F}(d) - \mathbb{E}[\hat{F}(d)]\right]^2. \end{aligned} \quad (10)$$

The first term is the variance of the modelled process and cannot be avoided. The second term measures the difference between the true mean of the process and the estimated mean. The third term is the variance of the estimated model around its mean. The above expression can be written as:

$$\text{Err}_{\mathcal{D}}(d) = \sigma_e^2 + \text{Bias}^2(\hat{F}(d)) + \text{Variance}(\hat{F}(d)) \quad (11)$$

The more complex we make the candidate transformation \hat{F} the lower the bias but higher the variance. Equation (11) expresses the transformation bias-variance tradeoff: simple transformations can only represent the training data’s coarse details (high bias) causing a high prediction error (low variance) because the transformation ignores important aspects of the data structure; complex transformations can represent training data structures in great detail (lower bias) but the prediction error increases (in variance) because the transformation do not generalise to other data.

The optimal transformation is the one that achieves the best generalization error on the new unseen samples. There are two types of methods to select the transformation that has the best generalization error: empirical methods use validation data different from the train data to assess the model generalization error on the test data, e.g., cross-validation and bootstrap; criteria based methods provide an estimate of the model generalization error on the test data based on the error on the training data and the complexity the model, e.g., Bayesian Information Criterion. The minimum description length criterion

is in the later group, and I chose it as the model selection criterion for feature space transformation.

4.2 The MDL Principle

Model selection is a widely studied subject, see [14], and the minimum description length (MDL) criterion is among the most common criteria of model selection. Rooted in information theory, the MDL principle was initially thought as a method to find the minimum number of bits required to transmit a particular message msg . To transmit this message a codebook cbk such as Huffman coding can be used to compress the message. Thus, the total number of bits required to transmit the message is

$$DL(msg, cbk) = DL(msg | cbk) + DL(cbk), \quad (12)$$

corresponding to the description length of the message msg encoded with the codebook cbk plus the description length of the codebook cbk . The MDL principle says that the optimal trade-off between these two quantities is achieved with the codebook cbk_{min} that minimizes the above expression. The minimum description length is written as

$$MDL(msg) = DL(msg | cbk_{min}) + DL(cbk_{min}), \quad (13)$$

where cbk_{min} is the optimal codebook that allows the message msg to be transmitted with the minimum number of bits.

The relation between the MDL criterion and the problem of model selection is straightforward: it assesses the trade-off between the data likelihood (the message) under a given model (the codebook) and the complexity of that model. In the problem we are addressing, the data \mathcal{D} will be transformed into a new feature-space by a transformation \hat{F} . Hence, Equation (12) is written as the sum of the likelihood of the data \mathcal{D} on the new feature space and the complexity of the feature-space transformation \hat{F} . Formally, we have

$$DL(\hat{F}, \mathcal{D}) = -\sum_{d \in \mathcal{D}} \log p(d | \hat{F}_i) + \frac{npars}{2} \cdot \log N, \quad (14)$$

where $npars$ is the number of parameters of the transformation \hat{F} , and N is the number of samples in the training dataset. Hence, the MDL criterion is designed “to achieve the best compromise between likelihood and ... complexity relative to the sample size”, [4]. Finally, the optimal feature-space transformation is the one

that minimizes Equation (14), which results in

$$F = \arg \min_{\hat{F}} \text{DL}(\hat{F}, \mathcal{D}). \quad (15)$$

The MDL criterion provides an estimate of the model error on the test data. Note that it is not an absolute estimate – it is only relative among candidate models. To evaluate the set Θ of candidate models and to better assess the characteristics of each model relatively to others we can compute the posterior probability of each model,

$$P(F_n | \mathcal{D}) = \frac{e^{-\frac{1}{2}\text{DL}(F_n)}}{\sum_{i=1}^{|\Theta|} e^{-\frac{1}{2}\text{DL}(F_i)}}. \quad (16)$$

The minimum description length approach is formally identical to the Bayesian Information Criterion but is motivated from a Bayesian perspective, see [25].

4.3 Dense Spaces Transformations

Some of the input feature spaces (depending on its media type) can be very dense making its modelling difficult due to cross-interference between classes. Expanding the original feature space into higher-dimensional ones results in a sparser feature space where the modelling of the data can be easier. This technique is applied by many related methods such as kernels. The discussion section of the next chapter will discuss these relationships.

The low-level visual features that I use are dense and low-dimensional: hence, keyword data may overlap thereby increasing the cross-interference. This means that not only the discrimination between keywords is difficult but also the estimation of a density model is less effective due to keyword data overlapping. One solution is to expand the original feature space into a higher-dimensional feature space where keywords data overlap is minimal. Thus, we define F_V as the transformation that increases the number of dimensions of a dense space with m dimensions into an optimal space with k_V dimensions

$$F_V(d_{V,1}, \dots, d_{V,m}) = \begin{bmatrix} f_{V,1}(d_{V,1}, \dots, d_{V,m}) \\ \vdots \\ f_{V,k_V}(d_{V,1}, \dots, d_{V,m}) \end{bmatrix}, \quad k_V \gg m. \quad (17)$$

In other words, for an input feature space with m dimensions the transformation $F_V(d_{V,1}, \dots, d_{V,m})$

generates a k_V dimensional feature space with $k_V \gg m$, where each dimension i of the new feature space corresponds to the function $f_{V,i}(d_{V,1}, \dots, d_{V,m})$. The optimal number of such functions will be selected by the MDL principle and the method to estimate the functions is defined next.

4.3.1 Visual Features Pre-Processing

The feature processing step normalises the features and creates smaller-dimensional subspaces from the original feature-spaces. The low-level visual features that we use in our implementation are:

- **Marginal HSV distribution moments:** this 12 dimensional colour feature captures the 4 central moments of each colour component distribution. I use 3 subspaces corresponding to the 3 colour components with 4 dimensions each subspace.
- **Gabor texture:** this 16 dimensional texture feature captures the frequency response (mean and variance) of a bank of filters at different scales and orientations. I use 8 subspaces corresponding to each filter response of 2 dimensions each.
- **Tamura texture:** this 3 dimensional texture feature is composed of the of image's coarseness, contrast and directionality.

I tiled the images in 3 by 3 parts before extracting the low-level features. This has two advantages: it adds some locality information and it greatly increases the amount of data used to learn the generic codebook.

4.3.2 Visual Transformation: Hierarchical EM

The original visual feature vector $d_V = (d_{V,1}, \dots, d_{V,m})$ is composed of several low-level visual features with a total of m dimensions. These m dimensions span the J visual feature types (e.g., marginal HSV colour moments, Gabor filters and Tamura), i.e. the sum of the number of dimensions of each one of the J visual feature space equals m . This implies that each visual feature type j is transformed individually by the corresponding $F_{V,j}(d_{V,j})$ and the output is concatenated into the vector

$$\mathbf{F}_V(d_V) = \begin{bmatrix} \mathbf{F}_{V,1}(d_{V,1}) \\ \vdots \\ \mathbf{F}_{V,j}(d_{V,j}) \end{bmatrix}, \quad (18)$$

where the dimensionality of the final \mathbf{F}_V transformation is the sum of the dimensionality of each individual visual feature space transformation $\mathbf{F}_{V,j}$, i.e.,

$$k_V = k_{V,1} + \dots + k_{V,j} + \dots + k_{V,J}. \quad (19)$$

The form of visual feature space transformations $\mathbf{F}_{V,j}$ is based on Gaussian mixture density models. The components of a GMM capture the different modes of the problem's data. I propose to use each component as a dimension of the optimal feature space where modes are split and well separated thereby creating a feature space where keywords can be modelled with a simple and low cost algorithm.

The transformations are defined under the assumption that subspaces are independent. This allows us to process each visual feature subspace j individually and model it as a Gaussian mixture model (GMM)

$$p(d_V) = p(d_V | \theta_j) = \sum_{m=1}^{k_{V,j}} \alpha_{m,j} p(d_V | \mu_{m,j}, \sigma_{m,j}^2), \quad (20)$$

where d_V is the low-level feature vector, θ_j represents the set of parameters of the model of the j visual feature subspace: the number $k_{V,j}$ of Gaussians components, the complete set of model parameters with means $\mu_{m,j}$, covariances $\sigma_{m,j}^2$, and component priors $\alpha_{m,j}$. The component priors have the convexity constraint $\alpha_{1,j}, \dots, \alpha_{k_{V,j},j} \geq 0$ and $\sum_{m=1}^{k_{V,j}} \alpha_{m,j} = 1$. Thus, for each visual feature space j , we have the Gaussian mixture model with $k_{V,j}$ components which now defines the transformation,

$$\mathbf{F}_{V,j}(d_V) = \begin{bmatrix} \alpha_{1,j} p(d_V | \mu_{1,j}, \sigma_{1,j}^2) \\ \vdots \\ \alpha_{k_{V,j},j} p(d_V | \mu_{k_{V,j},j}, \sigma_{k_{V,j},j}^2) \end{bmatrix}, \quad (21)$$

where each dimension corresponds to a component of the mixture model. The critical question that arises from the above expression is that one does not know the optimal complexity of the GMM in advance. The complexity is equivalent to the number of parameters, which in our case is proportional to

the number of mixture components $k_{V,j}$:

$$npars_j = k_{V,j} + \dim_j \cdot k_{V,j} + k_{V,j} \frac{\dim_j \cdot (\dim_j + 1)}{2}, \quad (22)$$

where \dim_j is the dimensionality of the visual subspace j . Note the relation between this equation and Equation (14). To address the problem of finding the ideal complexity we implemented a hierarchical EM algorithm that starts with a large number of components and progressively creates different GMM models with a decreasing number of components. For example, if it starts with 10 random components the EM will fit those 10 GMM components, store that model, deletes the weakest component and restarts the fitting with the previously 9 fitted components that will compensate the deleted component. The process is repeated until one component remains. In the end the algorithm generated 10 mixtures that are then assessed with the MDL criterion and the best one is selected. The implemented hierarchical EM adopts several other strategies that we will describe next.

4.3.2.1 Detailed Hierarchical EM

The hierarchical EM algorithm was implemented in C++ and it is based on the one proposed by Figueiredo and Jain [12]: it follows the component-wise EM algorithm with embedded component elimination.. The mixture fitting algorithm presents a series of strategies that avoids some of the EM algorithm’s drawbacks: sensitivity to initialization, possible convergence to the boundary of the parameter space and the estimation of different feature importance.

The algorithm starts with a number of components that is much larger than the real number and gradually eliminates the components that start to get few support data (singularities). This avoids the initialization problem of EM since the algorithm only produce mixtures with components that have enough support data. Component stability is checked by assessing its determinant (close to singularity) and its prior (few support data). If one of these two conditions is not met, we delete the component and continue with the remaining ones. This strategy can cause a problem when the initial number of components is too large: no component receives enough initial support causing the deletion of all components. To avoid this situation, component parameters are updated sequentially and not simultaneously as in standard EM. That is: first update component 1 parameters (μ_1, σ_1^2) , then

recompute all posteriors, update component 2 parameters (μ_2, σ_2^2) , recompute all posteriors, and so on.

After finding a good fit for a GMM with k components, the algorithm deletes the weakest component and restarts itself with $k - 1$ Gaussians and repeats the process until a minimum number of components is reached. Each fitted GMM is stored and in the end the set of fitted models describe the feature subspace at different levels of granularities.

The hierarchical EM algorithm for Gaussian mixture models addresses the objective of finding the optimal feature space by (1) creating transformations with different complexities and (2) splitting data modes into different space dimensions, hence enabling the application of low-cost keyword modelling algorithms.

4.3.3 Experiments

Experiments assessed the behaviour of the hierarchical EM algorithm on a real world photographic image collection. The collection is a 4,500 images subset of the widely used Corel CDs Stock Photos. More details regarding this collection are provided in Chapter 2. The visual features used in these experiments are the Gabor texture features, the Tamura texture features and the marginal HSV colour moments as described in Section 4.4.1.

The evolution of the model likelihood and complexity with a decreasing number of components are the two most important characteristics of the hierarchical EM that I wish to study. The algorithm is applied to individual visual feature subspaces. Each GMM model starts with $k_{V,j} = 200$ Gaussians, and the algorithm fits models with a decreasing number of components until a minimum number of Gaussians of 1.

One of the assumptions of the minimum description length principle is that the number of samples is infinite. Thus, to increase the accuracy of the MDL criterion we created 3 by 3 tiles of the training images. This increased the number of training samples by a factor of 9, which greatly improves the quality of the produced GMMs because of the existence of more data to support the model parameters.

4.3.4 Results and Discussion

An advantage of the chosen algorithm to find the optimal transformation is its natural ability to generate a series of transformations with different levels of complexities. This allows assessing different

GMMs with respect to the trade-off between decreasing levels of granularity and their fit to the data likelihood.

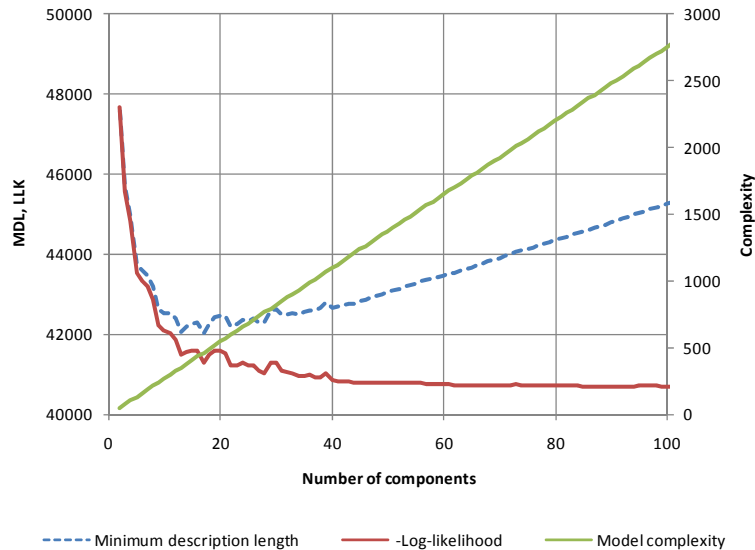


Figure 1. Model selection for the Gabor filters features (Core15000).

Figure 1 illustrates the output of a GMM model fitting to the output of one Gabor filter. The minimum description length curve (blue line) shows the trade-off between the models complexity (green line) and the models likelihood (red line). Note that we are actually plotting $-\log\text{-likelihood}$ for better visualization and comparison. The models likelihood curve is quite stable for models with a large number of components (above 40). On the other extreme of the curve one can see that models with fewer than 40 components the likelihood start to exhibit a poorer performance. The small glitches in the likelihood curve are the result of component deletion from a particularly good fit (more noticeable between 10 and 20 components). This effect is more visible when a component has been deleted from a model with a low number of components because the remaining ones are not enough to cover the data that was supporting the deleted one. The model complexity curve shows the penalty increasing linearly with the number of components according to Equation (22).

The most important curve of this graph is the minimum description length curve. At the beginning it closely follows the likelihood curve because the complexity cost is low. As the model complexity increases the model likelihood also becomes better but no longer at the same rate as initially (less than 10 components). This causes the model penalty to take a bigger part in the MDL formula, and after 20

components the MDL criterion indicates that those models are not better than previous ones. Thus, according to the MDL criterion the optimal transformation for this Gabor filter is the model with 18 components.




Figure 2. Model selection for the Tamura features (Corel5000).

The selection of the transformation of the Tamura visual texture features is illustrated in Figure 2. The behaviour is the same as for the Gabor features with the only difference that the change from the descending part of the MDL curve to the ascending part is not so pronounced. This indicates that the optimal model, $k_{V,j} = 39$, is not so distinct from the neighbouring models with $k_{V,j}$ between 30 and 50.

Finally, Figure 3 illustrates the optimal transformation selection experiments for a colour channel of the marginal HSV colour moments histograms. The behaviour is again similar to the previous ones and the optimal model, $k_{V,j} = 12$, is quite distinct from the surrounding neighbours. Note that the likelihood curve glitches are again present in this feature space which is an indication that the GMMs are well fitted to the data with a low number of components and that a deletion of a component leaves uncovered data causing the likelihood jitter.




Figure 3. Model selection for the marginal moments of HSV colour histogram features (Corel5000).

4.4 Sparse Spaces Transformations

Text features are high-dimensional sparse data, which pose some difficulties to parametric generative models because each parameter receives little data support. In discriminative models one observes over-fitting effects because the data representation might be too optimistic by leaving out a lot of the underlying data structure information. High-dimensional sparse data must be compressed into a lower dimensional space to ease the application of generative models. This optimal data representation is achieved with a transformation function defined as

$$F_T(d_{T,1}, \dots, d_{T,n}) = \begin{bmatrix} f_{T,1}(d_{T,1}, \dots, d_{T,n}) \\ \vdots \\ f_{T,k_T}(d_{T,1}, \dots, d_{T,n}) \end{bmatrix}, \quad k_T \ll n, \quad (23)$$

where n is the number of dimensions of the original sparse space, and k_T is the number of dimensions of the resulting optimal feature space.

In other words, the sparse spaces transformation $F_T(d_{T,1}, \dots, d_{T,n})$ receives as input a feature space with n dimensions and generates a k_T dimensional feature space, where each dimension i of the new optimal feature space corresponds to the function $f_{T,i}(d_{T,1}, \dots, d_{T,n})$. The optimal number of such functions will be selected by the MDL principle, and the method to estimate the functions is defined next.

4.4.1 Text Feature Pre-Processing

The text part of a document is represented by the feature vector $d_T = (d_{T,1}, \dots, d_{T,n})$ obtained from the text corpus of each document by applying several standard text processing techniques [41]: stop words are first removed to eliminate redundant information, and rare words are also removed to avoid over-fitting [19]. After this, the Porter stemmer [32] reduces words to their morphological root, which we call *term*. Finally, we discard the term sequence information and use a bag-of-words approach.

These text pre-processing techniques result in a feature vector $d_T = (d_{T,1}, \dots, d_{T,n})$, where each $d_{T,i}$ is the number of occurrences of term t_i in document d .

4.4.2 Text Codebook by Feature Selection

To reduce the number of dimensions in a sparse feature space we rank terms t_1, \dots, t_n by their importance to the modelling task and select the most important ones. The information gain criterion ranks the text terms by their importance, and the number of text terms is selected by the minimum description length. The criterion to rank the terms is the average mutual information technique, also referred to as information gain [41], expressed as

$$\text{IG}(t_i) = \frac{1}{L} \sum_{j=1}^L \text{MU}(y_j, t_i), \quad (24)$$

where t_i is term i , and y_j indicates the presence of keyword w_j . The information gain criterion is the average of the mutual information between each term and all keywords. Thus, one can see it as the mutual information between a term t_i and the keyword vocabulary.

The mutual information criterion assess the common entropy between a keyword entropy $H(y_j)$ and the keyword entropy given a term t_i , $H(y_j | t_i)$. Formally the mutual information criterion is defined as

$$\text{MU}(y_j, t_i) = \sum_{y_j \in \{0;1\}} \sum_{d_{T,i}} p(y_j, d_{T,i}) \log \frac{p(y_j, d_{T,i})}{p(y_j) p(d_{T,i})}, \quad (25)$$

where $d_{T,i}$ is the number of occurrences of term t_i in document d . Yang and Pedersen [44] and

Forman [13] have shown experimentally that this is one of the best criteria for feature selection. A document d is then represented by k_T text terms as the mixture

$$p(d) = \sum_{i=1}^{k_T} \alpha_i p(t_i | d) = \sum_{i=1}^{k_T} \alpha_i \frac{d_{T,i}}{|d|}, \quad (26)$$

where $d_{T,i}$ is number of occurrences of term t_i in document d . The parameters of the above mixture are the priors α_i of corresponding to term t_i . This results in a total number of parameters

$$npars = k_T. \quad (27)$$

A list of models is constructed by progressively adding terms to each model according to the order established by the information gain criterion. In this particular case of sparse text features the complexity of the transformation is equivalent to the number k_T of text terms. The application of the MDL criterion in Equation (14) is now straightforward.

Finally, terms are weighted by their inverse document frequency, resulting in the feature space transformation function

$$f_{T,i}(d_T) = -d_{T,r(i)} \cdot \log \left(\frac{N}{\text{DF}(d_{T,r(i)})} \right), \quad (28)$$

where N is the number of documents in the collection, $\text{DF}(d_{T,i})$ is the number of documents containing the term t_i , and $r(i)$ is a permutation function that returns the i^{th} text term of the information gain rank.

4.4.3 Experiments

Experiments assessed the behaviour of the information gain criterion on the Reuters news collection described in Chapter 2. The text corpus was processed as described in Section 4.5.1 to obtain the text terms, and models are constructed by adding terms to the model according to the information gain rank.

4.4.4 Results and Discussion

The evolution of the model likelihood and complexity with an increasing number of terms is again the most important characteristic that we wish to study. Figure 4 illustrates the model likelihood (red line)

versus the model complexity (green line) and the minimum description length criterion as a measure of their trade-off. Note that the graph is actually showing the $-\log\text{-likelihood}$ for easier visualization and comparison.

Figure 4 illustrates the improving likelihood as new terms are added to the feature space. The curve smoothness observed in this graph is due to the scale of the x-axis (100 times greater than in the images case) and to the fact that neighbouring terms have similar information value.

The problem of selecting the dimensionality of the optimal feature space is again answered by the minimum description length criterion that selects a feature space with 972 dimensions. It is interesting to notice that the MDL selects a low dimensionality reflecting a model with lower complexity than others with better likelihood but higher complexity. Note that if we had more samples (in this dataset the number of samples is limited to 7,770) we would be able to select a more complex model (remember that the MDL criterion assumes an infinite number of samples).

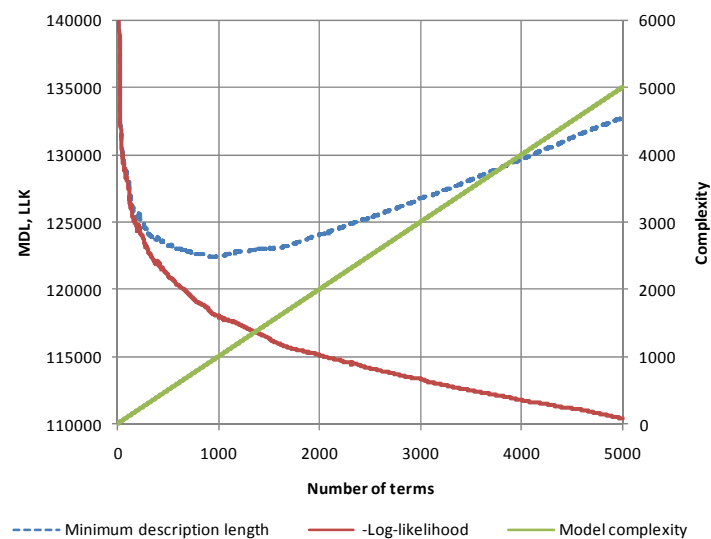


Figure 4. Model selection for the bag-of-word features (Reuters).

Moreover, information gain is a feature selection method that ranks terms by their discriminative characteristics and does not actually try to faithfully replicate the data characteristics. This is in contrast with the hierarchical EM method used for the dense feature spaces that is a pure generative approach. Hence, when adding new terms to the optimal feature space, we are directly affecting the classification performance.

5 Keyword Models

Modelling keywords in terms of multimedia information is the main objective of the first part of this thesis. Keywords are present in multimedia documents according to complex patterns that reflect their dependence and correlations. Different probability distributions can be applied to capture this information, also Bayesian networks can be used to define complex distributions that try to represent complex keyword interactions. This thesis opted to assume nothing about keyword interactions, and we define keywords as Bernoulli random variables with

$$p(y_t = 1) = 1 - p(y_t = 0) = \frac{|\mathcal{D}_{w_t}|}{|\mathcal{D}|}, \quad (29)$$

where y_t is a particular keyword, $|\mathcal{D}|$ is the size of the training collection and $|\mathcal{D}_{w_t}|$ is the number of documents in the training collection containing keyword w_t . In the previous chapter we proposed a probabilistic framework

$$p(y_t | F(d), \beta_t), \quad y_t = \{0, 1\}, \quad (30)$$

where $F(d)$ is a visual and text data transformation that creates a unique multi-modal feature space, and a keyword w_t is represented in that feature space by a model β_t . We will ignore the feature type and use a plain vector to represent the low-level features of a document as

$$F(d^j) = (F_T(d_T), F_V(d_V)) = (f_1^j, \dots, f_M^j). \quad (31)$$

One of the goals of the proposed $F(d)$ transformation is the creation of an optimal feature space, where simple and scalable keyword models β_t can be used. This chapter will propose the application of linear models to address this particular problem. The setting is a typical supervised learning problem, where documents are labelled with the keywords that are present in that document. Thus, we define

$$y^j = (y_1^j, \dots, y_L^j), \quad (32)$$

as the binary vector of keyword annotations of document j , where each y_t^j indicates the presence of keyword w_t in document j if $y_t^j = 1$. Note that a perfect classifier would have $(y - d_W) = 0$ on a

new document. The annotations vector y^j is used to estimate keyword models and to test the effectiveness of the computed models.

5.1 Keyword Baseline Models

The first linear models that we shall present in this section are simple but effective models that can be applied in the multi-modal feature space [26]. The advantage of both Rocchio classifier and naïve Bayes classifier is that they can be computed analytically.

5.1.1 Rocchio Classifier

Rocchio classifier was initially proposed as a relevance feedback algorithm to compute a query vector from a small set of positive and negative examples [34]. It can also be used for categorization tasks, e.g., [18]: a keyword w_t is represented as a vector β_t in the multi-modal space, and the closer a document is to this vector the higher is the similarity between the document and the keyword. A keyword vector β_t is computed as the average of the vectors of both relevant documents $\{\mathcal{D}_{w_t}\}$ and non-relevant documents $\{\mathcal{D} \setminus \mathcal{D}_{w_t}\}$,

$$\beta_t = \frac{1}{|\mathcal{D}_{w_t}|} \sum_{d \in \mathcal{D}_{w_t}} \frac{F(d)}{\|F(d)\|} - \frac{1}{|\mathcal{D} \setminus \mathcal{D}_{w_t}|} \sum_{d \in \mathcal{D} \setminus \mathcal{D}_{w_t}} \frac{F(d)}{\|F(d)\|}. \quad (33)$$

For retrieval scenarios, documents are ranked according to their proximity to the keyword vector. The cosine similarity measure has already proven to perform quite well in high-dimensional spaces:

$$\cos(\beta_t, F(d)) = \frac{\beta_t \cdot F(d)}{\|\beta_t\| \cdot \|F(d)\|} = \frac{\sum_{i=1}^T \beta_{t,i} \cdot f_i}{\sqrt{\sum_{i=1}^M (\beta_{t,i})^2} \cdot \sqrt{\sum_{i=1}^M (f_i)^2}} \quad (34)$$

The Rocchio classifier is a simple classifier that has been widely used in the area of text information retrieval and, as we have shown, can also be applied to semantic-multimedia information retrieval. Moreover, this classifier is particularly useful for online learning scenarios and other interactive applications where the models need to be updated on-the-fly or the number of training examples are limited.

5.1.2 Naïve Bayes Model

The naïve Bayes classifier assumes independence between feature dimensions and is the result of the direct application of Bayes's law to classification tasks:

$$p(y_t = 1 | d) = \frac{p(y_t = 1) p(d = f_1, \dots, f_M | y_t = 1)}{p(d)} \quad (35)$$

The assumption that features f_i are independent of each other in a document can be modelled by several different independent probability distributions. A distribution is chosen according to some constraints that we put on the independence assumptions. For example, if we assume that features f_i can be modelled as the simple presence or absence in a document then we consider a binomial distribution. If we assume that features f_i can be modelled as a discrete value to indicate the presence confidence in a document then we consider a multinomial distribution, see [28]. The binomial distribution over features f_i would be too limiting; the multinomial distribution over features f_i offers greater granularity to represent a feature value.

One can compute the log-odds and classify a document with the keywords that have a value greater than zero:

$$\log \frac{p(w_j = 1 | d)}{p(w_j = 0 | d)} = \log \frac{p(y_t = 1)}{p(y_t = 0)} + M \sum_{i=1}^M p(f_i | d) \log \frac{p(f_i | y_t = 1)}{p(f_i | y_t = 0)} \quad (36)$$

Formulating naïve Bayes in log-odds space has two advantages: it shows that naïve Bayes is a linear model and avoids decision thresholds in multi-categorization problems. In this case the keyword models become

$$\beta_{t,i} = \log \frac{p(f_i | y_t = 1)}{p(f_i | y_t = 0)}, \quad i = 1, \dots, M. \quad (37)$$

5.2 Keywords as Logistic Regression Models

Logistic regression is a statistical learning technique that has been applied to a great variety of fields, e.g., natural language processing [5], text classification [30], and image annotation [17]. In this section we employ a binomial logistic model to represent keywords in the multi-modal feature space. The expression of the binomial logistic regression is

$$p(y_t = 1 | \mathbf{F}(d), \beta_t) = \frac{1}{1 + \exp(\beta_t \cdot \mathbf{F}(d))} \quad (38)$$

The logistic regression model is also a linear model, which makes it a scalable and efficient solution for modelling keywords. The theory of Generalized Linear Models shows how to derive the logistic regression expression from a point of view of pure linear models and without making use of the log-odds as we did here. I shall develop this later in this chapter.

5.2.1 Regularization

As discussed by Nigan, Lafferty and McCallum [30] and Chen and Rosenfeld [8], logistic regression may suffer from over-fitting. This is usually because features are high-dimensional and sparse meaning that the regression coefficients can easily push the model density towards some particular training data points. Zhang and Oles [46] have also presented a study on the effect of different types of regularization on logistic regression. Their results indicate that with the adequate cost function (regularization), precision results are comparable to SVMs with the advantage of rendering a probabilistic density model.

An efficient and well known method of tackling over-fitting is to set a prior on the regression coefficients. As suggested by Nigan, Lafferty and McCallum [30] and Chen and Rosenfeld [8] I use a Gaussian prior \mathcal{N}_ξ for the regression coefficients,

$$\beta_* \sim \mathcal{N}_\xi(\mu_\xi, \sigma_\xi^2) \quad (39)$$

with mean $\mu_\xi = 0$ and σ_ξ^2 variance. The Gaussian prior imposes a cost on models β_* with large norms thus preventing optimization procedures from creating models that depend too much on a single feature space dimension. When introducing the Gaussian prior in the keyword model expression we obtain

$$p(y_t = 1 | d, \beta_t, \sigma_\xi^2) = p(y_t = 1 | d, \beta_t) p(\beta_t | \sigma_\xi^2), \quad (40)$$

which we will now use in the maximum likelihood estimation. We will drop the variance σ_ξ^2 of the Gaussian prior in our notation.

5.2.2 Maximum Likelihood Estimation

The log-likelihood function computes the sum of the log of the errors of each document in the collection \mathcal{D} :

$$l(\beta_t | \mathcal{D}) = \sum_{j \in \mathcal{D}} \log(p(y_t^j | F(d^j), \beta_t) p(\beta_t)) \quad (41)$$

For each keyword model the likelihood function tells us how well the model and those parameters represent the data. The model is estimated by finding the minimum of the likelihood function by taking the regression coefficients as variables:

$$\beta_t = \min_{\beta} l(\beta | \mathcal{D}) \quad (42)$$

For models where the solution can be found analytically, the computation of the regression coefficients is straightforward. In cases, where the analytical solution is not available typical numerical optimization algorithms are adequate.

The regression coefficients need to be found by a numerical optimization algorithm that iteratively approaches a solution corresponding to a local minimum of the log-likelihood function. To find the minimum of the log-likelihood function $l(\beta)$ with respect to β , I use the Newton-Raphson algorithm:

$$\beta^{new} = \beta^{old} - \left(\frac{\partial^2 l(\beta^{old})}{\partial \beta \partial \beta^T} \right)^{-1} \left(\frac{\partial l(\beta^{old})}{\partial \beta} \right) \quad (43)$$

The first-order derivative matrix is a vector with M elements corresponding to the dimension of the space resulting from the application of $F(d)$ to the original data. The second-order derivative, the Hessian matrix, is a square-matrix with $M \times M$ components. The Hessian matrix imposes a high computational complexity (both in time and space) on the parameter estimation algorithm. In multimedia information retrieval we use feature spaces with thousands of dimensions, meaning that the processing of the Hessian matrix is computationally too costly. For these reasons, we must use algorithms that are more suitable for such a large-scale problem.

5.2.3 Large-Scale Model Computation

When applying the Newton-Raphson algorithm to high-dimensional data the Hessian matrix often cannot be computed at a reasonable cost because it is too large and dense. Large scale Quasi-Newton methods are an adequate solution for our problem: instead of storing and computing the full Hessian matrix, these methods store a few vectors that represent approximations implicitly made in previous iterations of the algorithm. The L-BFGS algorithm (limited-memory Broyden-Fletcher-Goldfarb-Shanno)

is one of such algorithms, see [23] for details: “*The main idea of this method is to use curvature information from only the most recent iterations to construct the Hessian approximation. Curvature information from earlier iterations, which is less likely to be relevant to the actual behaviour of the Hessian at the current iteration, is discarded in the interest of saving storage.*”

The L-BFGS algorithm iteratively evaluates the log-likelihood function and its gradient, and updates the regression coefficients and the Hessian approximation. For the binomial logistic regression the log-likelihood function is

$$l(\beta_t) = \sum_{d^j \in \mathcal{D}} \left(y_t^j \beta_t F(d^j) - \log(1 + \exp(\beta_t F(d^j))) \right) - \lambda \beta_t^2, \quad (44)$$

$$\lambda = \frac{1}{2\sigma_\xi^2},$$

where for each example d^j the variable y_t^j is 1 if the example contains the keyword w_t and 0 otherwise. $F(d^j)$ is the nonlinear space transformation of the document features. To minimize the log-likelihood we need to use the gradient information to find the β_t where the log-likelihood gradient is zero, i.e.,

$$\frac{\partial l(\beta_t)}{\partial \beta_t} = 0 = \sum_{d^j \in \mathcal{D}} F(d^j) \left(y_t^j - p(y_t^j = 1 | \beta_t, F(d^j)) \right) - \lambda \beta_t. \quad (45)$$

These two last equations are the binomial logistic regression functions that the L-BFGS algorithm evaluates on each iteration to compute the β_t regression coefficients.

We use the FORTRAN implementation provided by Liu and Nocedal [24] to estimate the parameters of both linear logistic models and log-linear models. It has been shown that L-BFGS is the best optimization procedure for both maximum entropy [27] and conditional random fields models [35]. For more details on the limited-memory BFGS algorithm see [31].

6 Evaluation

The presented algorithms were evaluated with a retrieval setting on the Reuters-21578 collection, on a subset of the Corel Stock Photo CDs [10] and on a subset of the TRECVID2006 development data.

6.1 Collections

The three collections used in this evaluation are described in more detail in Chapter 2.

6.1.1 Reuters-21578

This is a widely used text dataset which allows comparing our results with others in the literature. Each document is composed by a text corpus, a title (which we ignore), and labelled categories. This dataset has several possible splits and we have used the *ModApte* split which contains 9,603 training documents and 3,299 test documents. This is the same evaluation setup used in several other experiments [19, 28, 30, 46]. Terms appearing less than 3 times were removed. Only labels with at least 1 document on the training set and the test set were considered leaving us with 90 labels. After these steps we ended with 7,770 labelled documents for training.

6.1.2 Corel Images

This dataset was compiled by Duygulu et al. [10] from a set of COREL Stock Photo CDs. The dataset has some visually similar concepts (jet, plane, Boeing), and some concepts have a limited number examples (10 or less). In their seminal paper, the authors acknowledge that fact and ignored the classes with these problems. In this paper we use the same setup as in [45], [7], [16], [21] and [11], which differs slightly from the one used in the dataset original paper, [10]. The retrieval evaluation scenario consists of a training set of 4,500 images and a test set of 500 images. Each image is annotated with 1-5 keywords from a vocabulary of 371 keywords. Only keywords with at least 2 images in the test set and training set each were evaluated, which reduced the number of vocabulary to 179 keywords. Retrieval lists have the same length as the test set, i.e. 500 items.

6.1.3 TRECVID

To test the similarity ranking on a multi-modal data we used the TRECVID2006 data: since only the training set is completely labelled, we randomly split the training English videos into 23,709 training documents and 12,054 test documents. We considered each document to be a key-frame plus the ASR text within a window of 6 seconds around that key-frame. Key-frames are annotated with the standard vocabulary of 39 keywords provided by NIST.

6.2 Experiment Design

To evaluate the proposed framework we deployed a retrieval experiment for all collections listed in the

previous section. The experiment methodology was as follows:

1. For a given algorithm and a given a multi-modal feature space
 - a. For each keyword in the considered collection
 - i. Estimate the keyword model on the training set by applying a cross-validation with 5 folds and 10 value iterations, as suggested in [20], to determine the ideal Gaussian prior variance σ_{ξ}^2
 - ii. Compute the relevance of each test document
 - iii. Rank all test documents by their relevance for the considered keyword
 - iv. Use the collection relevance judgments to measure the retrieval effectiveness of the considered rank
 - b. Repeat step a) for all keywords
 - c. Compute the mean average precision
2. Repeat for a different algorithm or a for different multi-modal feature space

The above methodology was repeated for all linear models that we presented in this chapter and for different multi-modal feature spaces. We considered the Reuters-21578 collection, the Corel5000 collection, the ASR part of the TRECVID2006, the key-frames of the TRECVID2006 and both key-frames and text of the TRECVID2006 development data, which makes a total of five collections.

6.3 Text-Only Models

The text-only models experiments on the Reuters-21578 collection evaluated the sparse data processing part of our framework. The optimal feature space was created with the average mutual information criterion as described in Chapter 4. All presented linear models were used in the evaluation.

6.3.1 Retrieval Effectiveness

Experiments in the Reuters dataset were evaluated with mean average precision, Figure 5, mean precision at 20, Figure 6, and precision-recall curves, Figure 7. All results were obtained with a 972 dimensional multi-modal feature space selected by the minimum description length criterion.




Figure 5. Reuters-21578 retrieval MAP evaluation.

When comparing the naïve Bayes model to the logistic regression model, results confirm what one would expect: naïve Bayes performs much worse than logistic regression (24.3% MAP versus 49.0% MAP). However, it is a surprise to see that Rocchio classifier is actually comparable to logistic regression – it obtained 49.7%. This supports the hypothesis that Reuters data is structured in a single cluster shape.




Figure 6. Reuters-21578 retrieval MP@20 evaluation.

Another reason why the Rocchio classifier performs so well on this dataset is that from all three classifiers it is the one that uses the simplest assumptions about data (organized as a high-dimensional sphere). The implications are that it is less prone to over-fit on classes with few training examples, unlike logistic regression. However, MP@20 values on Figure 5 show that logistic regression is actually more selective than Rocchio because it can do better on the top 20 retrieved documents: logistic regression obtained 39.3% while Rocchio obtained only 37.1%. precision-recall curves, Figure 7, offer a more detailed comparison of the models and confirm that logistic regression and Rocchio are very similar.




Figure 7. Precision-recall curve evaluation on the Reuters-21578.




Figure 8. Retrieval precision for different space dimensions (text-only models).

6.3.2 Model Complexity Analysis

We also studied the effect of the optimal space dimensionality by measuring the MAP on different spaces. The different multi-modal feature spaces were obtained by progressively adding new terms according to the average mutual information criterion. Figure 8 shows that after some number of terms (space dimension) precision do not increase because the information carried by the new terms is already present in the previous ones. The graph confirms that Rocchio is consistently better than logistic regression. Note that the MDL point (972 terms) achieves a good trade-off between the model complexity and the model retrieval effectiveness.

6.4 Image-Only Models

The image-only models experiment on the Corel Images collection evaluated the dense data processing part of the framework. The multi-modal feature space was created with the hierarchical EM algorithm described in Chapter 4. The different multi-modal feature spaces were obtained by

concatenating different colour and texture representations. As before, we evaluated all linear models that we presented in this chapter.

6.4.1 Retrieval Effectiveness

We first applied the MDL criterion to select a multi-modal feature space and then ran the retrieval experiments for all linear models. The space selected by the MDL criterion has 2,989 dimensions.

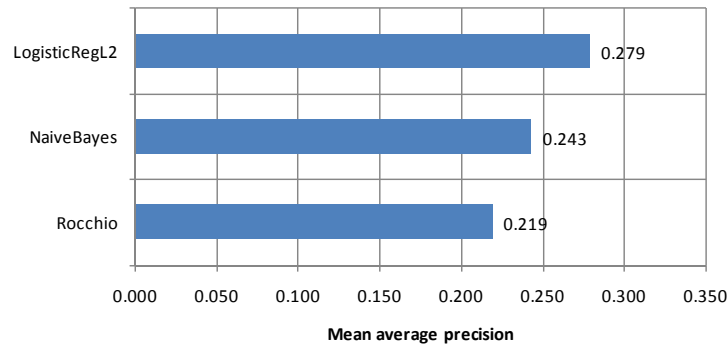


Figure 9. Corel retrieval MAP for different keyword models.

The MAP measures shown in Figure 9 shows that the best performance is achieved by the logistic regression models with a 27.9%, followed by naïve Bayes with 24.3% and Rocchio with 21.9%. The MP@20 measures in Figure 10 show that both naïve Bayes and logistic regression are affected similarly. However, the Rocchio classifier is less selective as the decrease in retrieval accuracy shows (from 21.9% to 10.1%). Contrary to the Reuters collection, the more complex structure of Corel Images dataset has affected the performance of the Rocchio classifier. Thus, both naïve Bayes and, more specifically, logistic regression can better capture the structure of this data.

The precision-recall curves in Figure 11 show that logistic regression is better than Rocchio and naïve Bayes across most of the recall area.

Results on this collection are more in agreement with what one would expect from the complexity of each model. Naïve Bayes applies a Gaussian on each dimension of the feature space, which reveals to be a more accurate assumption than the single cluster assumption made by the Rocchio classifier. Finally, logistic regression can better capture the non-Gaussian patterns of the data and achieve a better performance.




Figure 10. Corel retrieval MP@20 for different keyword models.




Figure 11. Precision-recall curves for different keyword models.

6.4.2 Model Complexity Analysis

Figure 12 depicts the evolution of the mean average precision with the dimensionality of the multi-modal feature space. Each point on the curve reflects the different levels of model complexities of the output of the hierarchical EM. Remember that the multi-modal feature space is the concatenation of the hierarchical EM Gaussian mixture models of the different feature subspaces. We concatenate sub-spaces with a similar number of level of complexity, e.g., GMMs with the same number of components per feature subspace.

For low dimensional multi-modal spaces the MAP for all models are quite low. Only when the dimensionality increases does the MAP achieve more stable values. The MAP stabilizes because the more complex GMMs models do not allow achieving a better discrimination between the relevant and non-relevant examples. The same phenomenon was observed on the Reuters collection.




Figure 12. Retrieval precision for different space dimensions.

6.5 Multi-Modal Models

For the multi-modal models we proceeded in the same way as for the other single-medium experiments with the difference that we deployed single-media and multi-modal experiments to compare and analyse the information value of each modality.

6.5.1 Retrieval Effectiveness

We first applied the MDL criterion to select a multi-modal feature space and then ran the retrieval experiments for all linear models. The space selected by the MDL criterion has 5,670 dimensions for the visual modality, 10,576 for the text modality, and the multi-modal space has a total of 16,247 dimensions. For the text modality the MDL selects the maximum number of terms because some of the key-frames have no ASR.

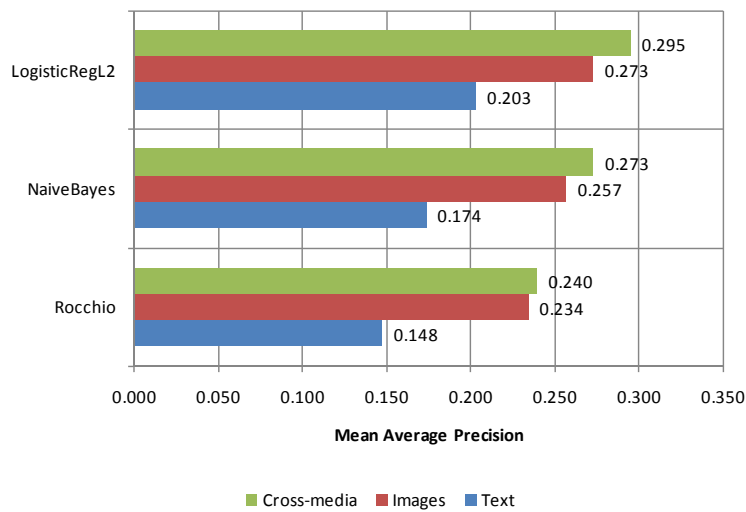


Figure 13. MAP by different modalities (TRECVID).




Figure 14. MP@20 by different modalities (TRECVID).

Figure 13 and Figure 14 present a summary of the retrieval effectiveness evaluation in terms of MAP and MP@20, respectively. All types of keyword models show the same variation with respect to each modality: text based models are always much lower than the image based models, and the difference between image based models and multi-modal models is always small. Moreover, logistic regression models are always better than naïve Bayes and Rocchio. This confirms previous knowledge that TRECVID collection is more difficult and its data exhibit a more complex structure, which is why logistic regression can exploit the non-Gaussian patterns of data: it achieves 20.2% MAP on the text-only experiment, 27.3% on the image-only experiment and 29.5% on the multi-modal experiment.




Figure 15. Precision-recall curve for the text-only models (TRECVID).

Text based models, Figure 15, exhibit a predictable behaviour: Rocchio is the less effective model, and

logistic regression is the most effective model for all values of recall. However, for values of recall higher than 70%, all models are very similar. Image based models, Figure 16, present a similar behaviour but the difference between the Rocchio and the naïve Bayes model is very small. It is also possible to observe that there is a significant difference between these two models for values of recall between 10% and 90%. Multi-modal models, Figure 17, show that naïve Bayes models better exploit the higher number of information sources than the Rocchio classifier. This is not a surprise as naïve Bayes considers individual dimensions, and the data structure is more complex than the spherical structure assumed by Rocchio. Also related to this phenomenon is the retrieval effectiveness obtained by the logistic regression model.

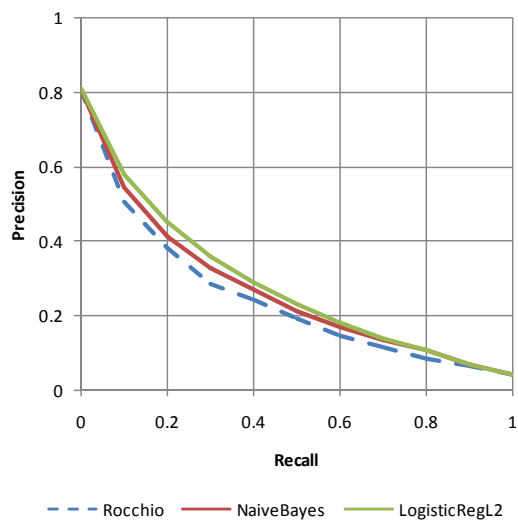


Figure 16: Precision-recall curve for image-only models (TRECVID).




Figure 17. Precision-recall curve for multi-modal models (TRECVID).

Finally, Figure 18 compares the logistic regression model on the different modalities. The first phenomenon to note is the difference between the text modality and the images modality. We believe that text-only models achieved such a low performance because some of the documents do not contain any text, and most concepts are more directly related to visual features than to text features. Multi-modal models perform better than the best single-media based models, which was a predictable behaviour given the increase in the number of predictors. However, this difference is not as big as we expected initially. We believe that the larger number of predictors would require a more exhaustive cross-validation procedure.

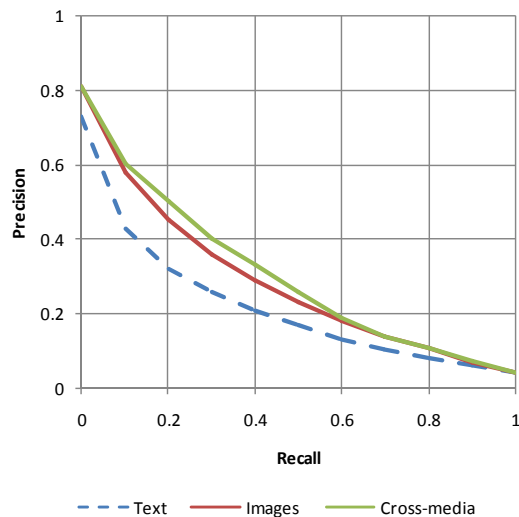


Figure 18. Precision-recall curves for different modalities.

6.5.2 Model Complexity Analysis

For the second experiment we studied the effect of the complexity of the feature space transformations – the number of dimensions of the optimal feature space. Figure 19 illustrates the text-based models’ retrieval effectiveness as new terms are added to the optimal feature space. The order by which terms are added is determined by the average mutual information. Retrieval effectiveness improves constantly but at a slower rate and with a different trend than for the Reuters collection. Again, we believe that this is related to the fact that some documents have no text and that TRECVID data is more complex.

Image based models, Figure 20, show an identical trend to the Corel collection. For a small number of dimensions the retrieval effectiveness is quite low and it quickly increases until a given dimensionality. The MAP achieves a stable range of values after around 5,000 dimensions and is not affected by the

addition of new dimensions to the feature space.

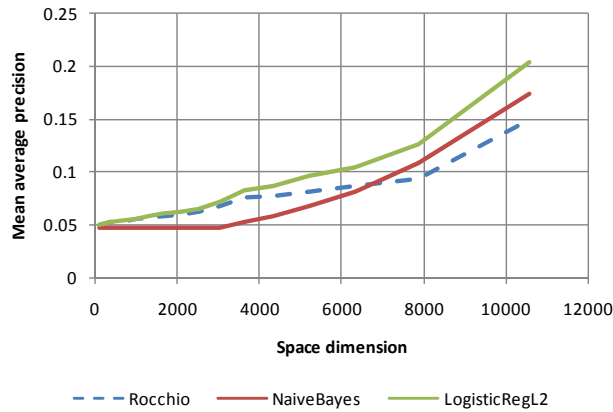


Figure 19. Retrieval precision for different space dimensions (TRECVID, text-only).

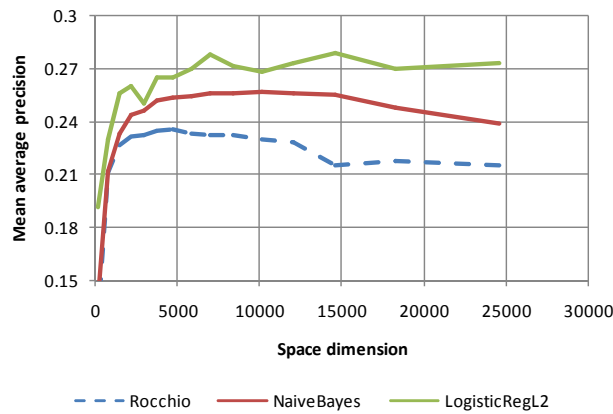


Figure 20. Retrieval precision for different space dimensions (TRECVID, image-only).




Figure 21. Retrieval precision for different space dimensions (TRECVID, multi-modal).

Multi-modal based models, Figure 21, exhibit a more irregular trend than the single-media models. The higher dimensionality and features heterogeneity might be the cause for this phenomenon. The

differences between the three models is related to the respective modelling capabilities: Rocchio assumes a spherical structure which reveals to be too simplistic for this data; naïve Bayes assumed independent dimensions, which is also not the best model for this data; finally, logistic regression further exploits feature dimensions interactions with linear combinations of them. Logistic regression, with an adequate cross-validation procedure, revealed to achieve the best retrieval effectiveness.

7 Conclusions

The creation of the multi-modal feature space is a generalization procedure which results in a trade-off between accuracy and computational complexity. Thus, the described algorithm offers an appealing solution for applications that require an information extraction algorithm with good precision, scalability, flexibility and robustness.

The novelty of the proposed framework resides in the simplicity of the linear combination of the heterogeneous sources of information that were selected by the minimum description length criterion.

7.1 Retrieval Effectiveness

The performed experiments show that our framework offers a performance in the same range as other state-of-the-art algorithms. Text and image results are quite good while multimodal experiments were affected by the noise present on the speech text and by the higher number of parameters to estimate. It was not surprising to see that logistic regression attains better results than naïve Bayes at the expense of a higher learning cost.

7.2 Model Selection

The algorithm's immunity to over-fitting is illustrated by the MAP curve stability as the model complexity increases. Logistic regression can be interpreted as ensemble methods (additive models) if we consider each dimension as a weak learner and the final model as a linear combination of those weak learners. This means that our model has some of the characteristics of additive models, namely the observed immunity to overfitting. It is interesting to note that the simple naïve Bayes model appears to be more immune to overfitting than the logistic regression model. This occurs because the optimization procedure fits the model tightly to the training data favouring large regression coefficients, while the naïve Bayes avoids overfitting by computing the weighted average of all codewords (dimensions). Note that

when fitting the model we are minimizing a measure of the model log-likelihood (the average classification residual error) and not a measure of how documents are ranked in a list (average precision). The mean average precision is the mean of the accumulated precision over a ranked list. Thus, we believe that if we trained our models with average precision as our goal metric, the retrieval results on the test set would improve.

7.3 Computational Scalability

Since the optimal feature space is common to all keywords the transformation must be computed only once for all keywords. Thus, the resources required to evaluate the relevancy of a multimedia document for each keyword are relatively small. During classification, both time and space complexity of the data representation algorithms is given by the number of Gaussians (clusters) selected by the model selection criteria. The computational complexity of linear models during the classification phase is negligible, resulting in a very low computational complexity for annotating multimedia content and making it quickly searchable.

The computational complexity during the learning phase is dominated by the hierarchical EM algorithm of mixture of Gaussians and the cross-validation method. The worst-case space complexity during learning is proportional to the maximum number of clusters, the number of samples, the dimension of each feature, and the total number of cross-validation iterations and folds. I consider this cost to be less important because the learning can be done offline.

Apart from the mixture of hierarchies [7] all other methods are some sort of kernel density distributions. It is well known [14] that the nature of these methods makes the task of running these models on new data computationally demanding: the model corresponds to the entire training set meaning that the demand on CPU time and memory increases with the training data.

Results show that such a low complexity approach compares competitively with much more complex approaches. It has a bearing on the design of image search engines, where scalability and response time is as much of a factor as the actual mean average precision of the returned results. In Chapter 7 I will illustrate how the low computational complexity of the proposed framework allows the development of new search paradigms that extract the semantic information of multimedia content on-the-fly.

7.4 Semantic Scalability

Assuming that the used set of keywords is a faithful sample of a larger keyword vocabulary it is expected that one can use the same optimal feature space to learn the linear model of new keywords and preserve the same models. Note that the optimal feature space is a representation of the data feature space: it is selected based on the entire data and independently of the number keywords. The consequence of this design is that systems can be semantically scalable in the sense that new keywords can be added to the system without affecting previous annotations.

8 References

- [1] A. Amir, J. O. Argillander, M. Campbell, A. Haubold, G. Iyengar, S. Ebabdollahi, F. Kang, M. Naphade, A. Natsev, J. R. Smith, J. Tesic, and T. Volkmer, "IBM Research TRECVID-2005 video retrieval system," in *TREC Video Retrieval Evaluation Workshop*, Gaithersburg, MD, USA, 2005.
- [2] J. Argillander, G. Iyengar, and H. Nock, "Semantic annotation of multimedia using maximum entropy models," in *IEEE Int'l Conf. on Acoustics, Speech, and Signal Processing*, Philadelphia, PA, 2005.
- [3] K. Barnard and D. A. Forsyth, "Learning the semantics of words and pictures," in *Int'l Conf. on Computer Vision*, Vancouver, Canada, 2001, pp. 408-415.
- [4] A. Barron and T. Cover, "Minimum complexity density estimation," *IEEE Transactions on Information Theory*, vol. 37, pp. 1034-1054, July 1991.
- [5] A. Berger, S. Pietra, and V. Pietra, "A maximum entropy approach to natural language processing," in *Computational Linguistics*, 1996, pp. 39-71.
- [6] D. Blei and M. Jordan, "Modeling annotated data," in *ACM SIGIR Conf. on research and development in information retrieval*, Toronto, Canada, 2003, pp. 127-134.
- [7] G. Carneiro and N. Vasconcelos, "Formulating semantic image annotation as a supervised learning problem," in *IEEE Conf. on Computer Vision and Pattern Recognition*, San Diego, CA, USA, 2005.
- [8] S. F. Chen and R. Rosenfeld, "A Gaussian prior for smoothing maximum entropy models," Carnegie Mellon University, Pittsburg, PA February 1999.
- [9] T. M. Cover and J. A. Thomas, *Elements of information theory*: John Wiley & Sons, 1991.
- [10] P. Duygulu, K. Barnard, N. de Freitas, and D. Forsyth, "Object recognition as machine translation: Learning a lexicon for a fixed image vocabulary," in *European Conf. on Computer Vision*, Copenhagen, Denmark, 2002, pp. 97-112.
- [11] S. L. Feng, V. Lavrenko, and R. Manmatha, "Multiple Bernoulli relevance models for image and video annotation," in *IEEE Conf. on Computer Vision and Pattern Recognition*, Cambridge, UK, 2004, pp. 1002-1009.
- [12] M. Figueiredo and A. K. Jain, "Unsupervised learning of finite mixture models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, pp. 381-396, March 2002.
- [13] G. Forman, "An extensive empirical study of feature selection metrics for text classification," *Machine Learning Research*, pp. 1289-1305, March 2003.
- [14] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: Data mining, inference and prediction*. Springer, 2001.
- [15] T. Hofmann, "Probabilistic latent semantic indexing," in *ACM SIGIR Conf. on research and development in information retrieval*, Berkeley, CA, USA, 1999.
- [16] J. Jeon, V. Lavrenko, and R. Manmatha, "Automatic image annotation and retrieval using cross-media relevance models," in *ACM SIGIR Conf. on research and development in information retrieval*, Toronto, Canada, 2003, pp. 119-126.

- [17] J. Jeon and R. Manmatha, "Using maximum entropy for automatic image annotation," in *Int'l Conf on Image and Video Retrieval*, Dublin, Ireland, 2004, pp. 24-32.
- [18] T. Joachims, "A probabilistic analysis of the Rocchio algorithm with TFIDF for text categorization," in *Int'l Conf. on Machine Learning*, Nashville, US, 1997.
- [19] T. Joachims, "Text categorization with Support Vector Machines: learning with many relevant features," in *European Conf. on Machine Learning*, 1998.
- [20] R. Kohavi, "A study of cross-validation and bootstrap for accuracy estimation and model selection," in *International joint Conference on artificial intelligence* Montréal, Québec, Canada, 1995.
- [21] V. Lavrenko, R. Manmatha, and J. Jeon, "A model for learning the semantics of pictures," in *Neural Information Processing System Conf.*, Vancouver, Canada, 2003.
- [22] S. Lazebnik, C. Schmid, and J. Ponce, "A maximum entropy framework for part-based texture and object recognition," in *Int'l Conf. on Computer Vision*, Beijing, China, 2005.
- [23] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," *Mathematical Programming*.
- [24] D. C. Liu and J. Nocedal, "On the limited memory method for large scale optimization," *Mathematical Programming B*, vol. 45, pp. 503-528.
- [25] D. J. C. MacKay, *Information theory, inference, and learning algorithms*. Cambridge: Cambridge University Press, 2004.
- [26] J. Magalhães and S. Rüger, "High-dimensional visual vocabularies for image retrieval," in *ACM SIGIR Conf. on research and development in information retrieval*, Amsterdam, The Netherlands, 2007.
- [27] R. Malouf, "A comparison of algorithms for maximum entropy parameter estimation," in *Sixth Conf. on Natural Language Learning*, Taipei, Taiwan, 2002, pp. 49-55.
- [28] A. McCallum and K. Nigam, "A comparison of event models for naive Bayes text classification," in *AAAI Workshop on Learning for Text Categorization*, 1998.
- [29] M. R. Naphade and T. S. Huang, "A probabilistic framework for semantic video indexing filtering and retrieval," *IEEE Transactions on Multimedia*, vol. 3, pp. 141-151, March 2001.
- [30] K. Nigam, J. Lafferty, and A. McCallum, "Using maximum entropy for text classification," in *IJCAI - Workshop on Machine Learning for Information Filtering*, Stockholm, Sweden, 1999, pp. 61-67.
- [31] J. Nocedal and S. J. Wright, *Numerical optimization*. New York: Springer-Verlag, 1999.
- [32] M. F. Porter, "An algorithm for suffix stripping," *Program*, vol. 14, pp. 130-137.
- [33] J. Rissanen, "Modeling by shortest data description," *Automatica*, vol. 14, pp. 465-471, 1978.
- [34] J. Rocchio, "Relevance feedback in information retrieval," in *The SMART Retrieval System: Experiments in Automatic Text Retrieval*, G. Salton, Ed.: Prentice-Hall, 1971, pp. 313-323.
- [35] F. Sha and F. Pereira, "Shallow parsing with conditional random fields," in *Human Language Technology Conf. of the North American Chapter of the Association for Computational Linguistics*, Edmonton, Canada, 2003.
- [36] C. G. M. Snoek, J. C. v. Gemert, T. Gevers, B. Huurnink, D. C. Koelma, M. v. Liempt, O. d. Rooij, K. E. A. v. d. Sande, F. J. Seinstra, A. W. M. Smeulders, A. H. C. Thean, C. J. Veenman, and M. Worring, "The MediaMill TRECVID 2006 semantic video search engine," in *TREC Video Retrieval Evaluation Workshop*, Gaithersburg, MD, USA, 2006.
- [37] C. G. M. Snoek, M. Worring, J.-M. Geusebroek, D. C. Koelma, F. J. Seinstra, and A. W. M. Smeulders, "The semantic pathfinder: using an authoring metaphor for generic multimedia indexing," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, pp. 1678-1689, October 2006.
- [38] A. Vailaya, M. Figueiredo, A. K. Jain, and H. J. Zhang, "Image classification for content-based indexing," *IEEE Transactions on Image Processing*, vol. 10, pp. 117-130, January 2001.
- [39] T. Westerveld and A. P. de Vries, "Experimental result analysis for a generative probabilistic image retrieval model," in *ACM SIGIR Conf. on research and development in information retrieval*, Toronto, Canada, 2003, pp. 135-142.
- [40] T. Westerveld, A. P. de Vries, T. Ianeva, L. Boldareva, and D. Hiemstra, "Combining information sources for video retrieval," in *TREC Video Retrieval Evaluation Workshop*, Gaithersburg, MD, USA, 2003.
- [41] Y. Yang, "An evaluation of statistical approaches to text categorization," *Information Retrieval*, pp. 69-90.

- [42] Y. Yang and C. G. Chute, "An example-based mapping method for text categorization and retrieval," *ACM Transactions on Information Systems*, vol. 13, pp. 252-277, July 1994.
- [43] Y. Yang and X. Liu, "A re-examination of text categorization methods," in *SIGIR*, 1999.
- [44] Y. Yang and J. O. Pedersen, "A comparative study on feature selection in text categorization," in *Int'l Conf. on Machine Learning*, Nashville, Tennessee, USA, 1997, pp. 412-420.
- [45] A. Yavlinsky, E. Schofield, and S. Rüger, "Automated image annotation using global features and robust nonparametric density estimation," in *Int'l Conf. on Image and Video Retrieval*, Singapore, 2005.
- [46] T. Zhang and F. J. Oles, "Text categorization based on regularized linear classification methods," *Information Retrieval*, pp. 5-31, 2001.