PROCEEDINGS

DATA COMPRESSION

CONFERENCE

March 30 - April 2, 1993 • Snowbird, Utah

Edited by
James A. Storer
Martin Cohn

Sponsored by the IEEE Computer Society
Technical Committee on Computer Communications

IEEE COMPUTER SOCIETY PRESS

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

DCC '93
Data Compression Conference

Edited by
James A. Storer
Martin Cohn

Sponsored by the **IEEE Computer Society**
Technical Committee on Computer Communications
in cooperation with **NASA/CESDIS**

IEEE Computer Society Press
Los Alamitos, California

Washington • Brussels • Tokyo

Published by the
IEEE Computer Society Press
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264

IEEE Computer Society Press Order Number 3392-02
IEEE Catalog Number 93TH0536-3
ISBN 0-8186-3391-3 (microfiche)
ISBN 0-8186-3392-1 (case)
ISSN 1068-0314

Additional copies can be ordered from

IEEE Computer Society Press
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1264

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331

IEEE Computer Society
13, avenue de l'Aquilon
B-1200 Brussels
BELGIUM

IEEE Computer Society
Ooshima Building
2-19-1 Minami-Aoyama
Minato-ku, Tokyo 107
JAPAN

Editorial production by Phyllis A. Walker and Mary E. Kavanaugh
Cover by Alex Torres
Printed in the United States of America by Braun-Brumfield, Inc.

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC.

DATA COMPRESSION CONFERENCE

1993

DCC 1993

This book contains the presentations from the third Data Compression Conference held March 30 - April 2, 1993, at Cliff Lodge, Snowbird, Utah. This conference was sponsored by the IEEE Computer Society Technical Committee on Computer Communications (TCCC) in cooperation with NASA/CESDIS.

General Chair: James A. Storer, *Brandeis U.*
Program Chair: Martin Cohn, *Brandeis U.*
Publicity Chair: Raymond Miller, *NASA/CESDIS*

Program Committee

Alberto Apostolico, *Purdue U./U. Padova*
Ronald B. Arps, *IBM*
Richard L. Baker, *PictureTel Inc.*
Anselm Blumer, *Tufts U.*
Abraham Bookstein, *U. Chicago*
Martin Cohn, *Brandeis U.*
Thomas M. Cover, *Stanford U.*
Robert M. Gray, *Stanford U.*
Daniel S. Hirschberg, *U.C. Irvine*
Arnaud Jacquin, *AT&T Labs*
Glen Langdon, *U.C. Santa Cruz*
Abraham Lempel, *Technion*
Thomas Lookabough, *Com. Labs*
Bradley J. Lucier, *Purdue U.*
Victor Miller, *IBM*
John H. Reif, *Duke U.*
James A. Storer, *Brandeis U.*
James C. Tilton, *NASA*
Jeffrey S. Vitter, *Brown U.*
Gregory K. Wallace, *DEC*
Victor K. Wei, *Bellcore*
Ian H. Witten, *U. Waikato*
Jacob Ziv, *Technion*

Acknowledgments: *The Program Committee gratefully acknowledges the support of NASA/CESDIS and the Computer Science Department of Brandeis University. We thank Nancy Campbell, Georgia Flanagan, and Anne Marie Murphy of NASA/CESDIS and Myrna Fox and Cathy Rossi of Brandeis University for their hard work on the conference administration.*

Offices of the IEEE Computer Society

Headquarters Office
1730 Massachusetts Avenue, N.W.
Washington, DC 20036-1903
Phone: (202) 371-0101 — Fax: (202) 728-9614

Publications Office
P.O. Box 3014
10662 Los Vaqueros Circle
Los Alamitos, CA 90720-1264
Membership and General Information: (714) 821-8380
Publication Orders: (800) 272-6657 — Fax: (714) 821-4010

European Office
13, avenue de l'Aquilon
B-1200 Brussels, BELGIUM
Phone: 32-2-770-21-98 — Fax: 32-3-770-85-05

Asian Office
Ooshima Building
2-19-1 Minami-Aoyama, Minato-ku
Tokyo 107, JAPAN
Phone: 81-3-408-3118 — Fax: 81-3-408-3553

IEEE Computer Society

IEEE Computer Society Press Publications

Monographs: A monograph is an authored book consisting of 100-percent original material.
Tutorials: A tutorial is a collection of original materials prepared by the editors, and reprints of the best articles published in a subject area. Tutorials must contain at least five percent of original material (although we recommend 15 to 20 percent of original material).
Reprint collections: A reprint collection contains reprints (divided into sections) with a preface, table of contents, and section introductions discussing the reprints and why they were selected. Collections contain less than five percent of original material.
Technology series: Each technology series is a brief reprint collection — approximately 126-136 pages and containing 12 to 13 papers, each paper focusing on a subset of a specific discipline, such as networks, architecture, software, or robotics.
Submission of proposals: For guidelines on preparing CS Press books, write the Editorial Director, IEEE Computer Society Press, P.O. Box 3014, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1264, or telephone (714) 821-8380.

Purpose

The IEEE Computer Society advances the theory and practice of computer science and engineering, promotes the exchange of technical information among 100,000 members worldwide, and provides a wide range of services to members and nonmembers.

Membership

All members receive the acclaimed monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others seriously interested in the computer field.

Publications and Activities

***Computer* magazine:** An authoritative, easy-to-read magazine containing tutorials and in-depth articles on topics across the computer field, plus news, conference reports, book reviews, calendars, calls for papers, interviews, and new products.

Periodicals: The society publishes six magazines and five research transactions. For more details, refer to our membership application or request information as noted above.

Conference proceedings, tutorial texts, and standards documents: The IEEE Computer Society Press publishes more than 100 titles every year.

Standards working groups: Over 100 of these groups produce IEEE standards used throughout the industrial world.

Technical committees: Over 30 TCs publish newsletters, provide interaction with peers in specialty areas, and directly influence standards, conferences, and education.

Conferences/Education: The society holds about 100 conferences each year and sponsors many educational activities, including computing science accreditation.

Chapters: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.

Technical Sessions

Session 1

Session 2

Session 3

Session 4

Session 5

Session 6

Session 7

Session 8

Session 9

Session 10

Posters Session

Technical Sessions

Combining Image Classification and Image Compression Using Vector Quantization

Karen L. Oehler Robert M. Gray

Information Systems Lab, Department of Electrical Engineering
Stanford University, Stanford, CA 94305-4055

Abstract

We describe a technique for combining compression and low-level classification of images. The goal is to produce codes with implicit classification information, where the compressed image incorporates classification information without further signal processing. This technique can provide direct low level classification or provide an efficient front end to more sophisticated full-frame recognition algorithms. Vector quantization (VQ) is a natural choice because two of its design components, clustering and tree-structured classification methods, have obvious applications to the pure classification problem as well as to the compression problem. Here we explicitly incorporate a Bayes risk component into the distortion measure used for code design in order to permit a tradeoff of mean squared error with classification error. This method is used to analyze simulated data, identify tumors in computerized tomography (CT) lung images, and identify man-made regions in aerial images.

1 Introduction

Image compression and classification play very important roles in making digital images useful. Because lossless compression (compression without error) can typically provide a compression ratio of only 2 - 4, lossy compression techniques become necessary for further compression. VQ is a common method of lossy compression that applies statistical techniques to optimize distortion/bit rate tradeoffs [1, 2, 3].

In many scientific and medical applications, images are used by human experts to make specific decisions or inferences. The actions of human experts can sometimes be mimicked by sophisticated computer algorithms, typically requiring large amounts of computer processing on entire images in order to reach decisions. Simple low-level local classification involving only small regions of an image, however, can assist human observers by highlighting special areas of interest and can simplify future classification/analysis algorithms by incorporating pre-processing into the digital representation. The approach described here is applicable to classification tasks where the goal is to classify fairly small regions in an image, for example, to identify

0-8186-3392-1/93 $3.00 © 1993 IEEE

lung nodules in a CT lung scan, microcalcifications in mammograms, or manmade objects in an aerial photograph. Our goal is to combine such local classification with the compression process to obtain a single code that does a good job of both, with an adjustable relative importance weighting of the two aspects. VQ provides a natural means to this end because two of its design components, clustering and tree-structured classification methods, have obvious applications to the pure classification problem as well as to the compression problem.

2 Vector Quantization

Vector quantization operates on individual image subblocks or vectors (e.g. blocks of 2 pixels $\times$ 2 pixels). For each vector, the VQ encoder determines the nearest codeword and outputs the chosen codeword's index. The sequence of indices so generated can then be stored or transmitted. The VQ decoder reverses this process as it inputs each index and outputs the appropriate codeword by simple table lookup. A clear advantage of VQ is the low complexity decoding process. As the input vectors mapping into a particular codeword will generally have much in common, they are likely to belong to the same class. By attaching class significance to the codewords, we can obtain both classification and compression with a single encoding step.

The accuracy of the classifier is measured by the Bayes risk, which reduces to a simple error probability if unit costs are chosen. The more general form allowing varying costs for different error types is useful for many applications (such as classifying tissue as tumor or healthy) and our algorithm provides a convenient means of including these costs. This Bayes risk term can be incorporated into the distortion measure, and the resulting codebook is then designed to determine class membership when the image is decompressed. Our method explicitly incorporates the performance of optimal (Bayes) classifiers operating on the vector quantized data, both through the form of the distributions (which are learned by training and do not require an a priori model), and through the explicit incorporation of a Bayes risk term into the average distortion that is minimized by the design algorithm. The extra complexity is almost entirely in the code design phase; the implemented code is only slightly more complicated than an ordinary vector quantizer. The classification requires no more bits to describe than the bits required for compression alone, an important feature in low memory or low bandwidth situations. As the classification information is entirely contained in the selected codeword, it can be easily displayed with the decompressed image. For example, a physician could select either a monochrome decompressed image or one with all the suspected anomalies colored pink for highlighting.

There have been several papers devoted to the related topic of designing a classifier that uses a VQ encoder. Kohonen [4, 5, 6] proposed a variety of "likelihood vector quantizers" or LVQ to perform classification using a VQ encoder and codebook, where the encoder operates as an ordinary minimum mean squared error selection of a representative from the codebook, but the codebook is designed in a manner that attempts to reduce classification error implicitly rather than reducing mean squared error. His general goal is to imitate a Bayes classifier with less complexity than other neural network approaches, but there is no explicit minimization of Bayes risk in the code design. The Kohonen approach differs from ours in many ways. First, our approach considers mean squared error in the optimization and his does not. As we allow a weighting of relative importance, his goal can be considered as an extreme point of ours. Compression ability is not explicitly considered in LVQ. His encoder uses only minimum mean squared error, even within the learning set, while we explicitly apply an empirical Bayes classifier as part of the encoding, balancing minimum squared error with empirical Bayes risk. We explicitly incorporate the Bayes risk into the average distortion measure being minimized, while Kohonen uses a heuristic to argue that moving centroids according to nearby class membership should asymptotically have the effect of approximating a Bayes risk.

3 Approach

Given a n-dimensional vector $X = (x_1, \ldots, x_n)^T$, a full search (unconstrained) VQ maps each input vector onto a set of codewords called the codebook: $q(X) = \hat{Y}_i$, where $\hat{Y}_i \in \{\hat{Y}_1, \ldots, \hat{Y}_N\}$. Denote the partition of input vector space induced by q by $\mathcal{P} = \{\mathcal{A}_1, \ldots, \mathcal{A}_N\}$ where $X \in \mathcal{A}_i$ if $q(X) = \hat{Y}_i$. Then $\alpha(X) = i$ if $X \in \mathcal{A}_i$ is the <u>encoder</u> and $\beta(i) = \hat{Y}_i$ is the <u>decoder</u> for the VQ q. Hence $q(X) = \beta(\alpha(X))$. Because there are many more possible input vectors than codewords, the quantizer compresses the data.

Suppose we have a collection of classes $\mathcal{H} = \{H_1, H_2, \ldots, H_M\}$. To perform a classification, we associate a class label with each cell $\mathcal{A}_i$ in the partition or, equivalently, with each index i produced by the encoder α, for $i \in \{1, \ldots, N\}$. The index can then be used to decompress the vector using the codebook and to classify the vector using the decision rule. Let $\delta(i) \in \mathcal{H}$, $i = 1, \ldots, N$, denote the decision function.

Codebook construction requires a labeled learning set $\mathcal{S} = \{X_1, \ldots, X_L\}$. Each training vector X_i is assigned to a specific class in $\mathcal{H}$. The samples provide the empirical distribution of X, used to train the vector quantizer.

The average ordinary distortion corresponding to (α, β) is

$$D(\alpha, \beta) = E[\, d(X, \beta(\alpha(X)))\,] \geq E[\min_i d(X, \beta(i))\,] \text{ where } 1 \leq i \leq N. \tag{1}$$

The mean squared error $d(X, \beta(\alpha(X))) = \|X - \beta(\alpha(X))\|^2$ is commonly used as the distortion measure.

The classification performance can be measured by the Bayes risk [7]:

$$B(\alpha, \delta) = \sum_{i=1}^{M} \sum_{j=1}^{M} P(\delta(\alpha(X)) = H_j | X \in H_i)\, P(H_i)\, C_{ij},$$

where C_{ij} is the relative cost of error incurred when $X \in H_i$ and $\delta(\alpha(X)) = H_j$. Often, $C_{ij} = 0$ when $i = j$ (correct decisions have zero cost). Note that the reproduction values supplied by β do not affect B. Define an indicator function $I_{H_i}(H_j)$ as 1 if