cludes these features initially for zero cost. It is also assumed that the feature acquisition (FA) cost associated with each feature is known in advance, and that the FA cost for a given feature is the same for all instances. Finally, CFA requires that the base-level classifiers produce not only a classification, but also a confidence (or posterior probability).

CFA trains an ensemble of classifiers $M_0 \ldots M_f$ that use successively larger subsets of the features to classify instances. M_0 uses only the "free" (zero cost) features, and M_1 additionally incorporates costly features F_1 through F_i. CFA reduces FA cost in that model M_i is trained only on instances that cannot be classified with sufficient confidence by model M_{i-1}. Therefore, values for feature F_i are acquired only for the instances that require it. At test time, each test instance is successively classified by M_0, M_1, $M_2 \ldots$ until its classification is sufficiently confident (i.e., until the confidence of the prediction reaches the confidence threshold). Again, features are acquired for the new instance only as required. In an empirical comparison with an existing method (Cost-Sensitive Naive Bayes) that makes acquisition decisions only during test time (and therefore requires that all training items be fully acquired), CFA achieves the same (or higher) level of performance at a much reduced cost (by at least an order of magnitude).

This work was done by Kiri L. Wagstaff of Caltech and Marie desJardins and James MacGlashan of the University of Maryland for NASA's Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov. NPO-46886

Algorithm for Lossless Compression of Calibrated Hyperspectral Imagery

NASA's Jet Propulsion Laboratory, Pasadena, California

A two-stage predictive method was developed for lossless compression of calibrated hyperspectral imagery. The first prediction stage uses a conventional linear predictor intended to exploit spatial and/or spectral dependencies in the data. The compressor tabulates counts of the past values of the difference between this initial prediction and the actual sample value. To form the ultimate predicted value, in the second stage, these counts are combined with an adaptively updated weight function intended to capture information about data regularities introduced by the calibration process. Finally, prediction residuals are losslessly encoded using adaptive arithmetic coding.

Algorithms of this type are commonly tested on a readily available collection of images from the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) hyperspectral imager. On the standard calibrated AVIRIS hyperspectral images that are most widely used for compression benchmarking, the new compressor provides more than 0.5 bits/sample improvement over the previous best compression results.

The algorithm has been implemented in Mathematica. The compression algorithm was demonstrated as beneficial on 12-bit calibrated AVIRIS images.

This work was done by Aaron B. Kiely and Matthew A. Klimesh of Caltech for NASA's Jet Propulsion Laboratory. For more information, contact iaoffice@jpl.nasa.gov. NPO-46547

Universal Decoder for PPM of any Order

Complexity can be reduced and flexibility increased, at small cost in performance.

NASA's Jet Propulsion Laboratory, Pasadena, California

A recently developed algorithm for demodulation and decoding of a pulse-position-modulation (PPM) signal is suitable as a basis for designing a single hardware decoding apparatus to be capable of handling any PPM order. Hence, this algorithm offers advantages of greater flexibility and lower cost, in comparison with prior such algorithms, which necessitate the use of a distinct hardware implementation for each PPM order. In addition, in comparison with the prior algorithms, the present algorithm entails less complexity in decoding at large orders.

An unavoidably lengthy presentation of background information, including definitions of terms, is prerequisite to a meaningful summary of this development. As an aid to understanding, the figure illustrates the relevant processes of coding, modulation, propagation, demodulation, and decoding. An M-ary PPM signal has M time slots per symbol period. A pulse (signifying 1) is transmitted during one of the time slots; no pulse (signifying 0) is transmitted during the other time slots.

The information intended to be con-

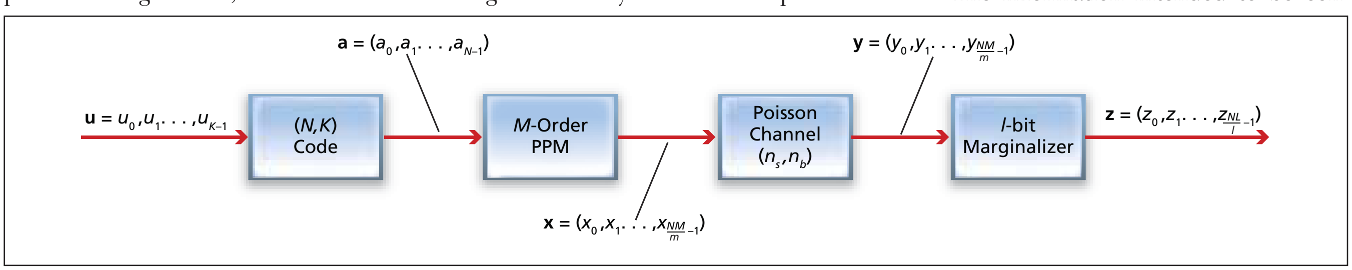

Processing of Information in an M-ary PPM communication system includes the sequence of steps depicted here. The l-bit marginalizer is a feature of the innovation reported here; the other features are typical of PPM systems in general.