

Two Information-Theoretic Tools to Assess the Performance of Multi-class Classifiers

Francisco J. Valverde-Albacete*, Carmen Peláez-Moreno

Departamento de Teoría de la Señal y de las Comunicaciones,
Universidad Carlos III de Madrid
Avda de la Universidad, 30. 28911 Leganés, Spain

Abstract

We develop two tools to analyze the behavior of multiple-class, or multi-class, classifiers by means of entropic measures on their confusion matrix or contingency table. First we obtain a balance equation on the entropies that captures interesting properties of the classifier. Second, by normalizing this balance equation we first obtain a 2-simplex in a three-dimensional entropy space and then the de Finetti entropy diagram or *entropy triangle*. We also give examples of the assessment of classifiers with these tools.

Key words: Multiclass classifier, confusion matrix, contingency table, performance measure, evaluation criterion, de Finetti diagram, entropy triangle

1. Introduction

Let $V_X = \{x_i\}_{i=1}^n$ and $V_Y = \{y_j\}_{j=1}^p$ be sets of input and output class identifiers, respectively, in a multiple-class classification task. The basic classification event consists in “presenting a pattern of input class x_i to the classifier to obtain output class identifier y_j ,” ($X = x_i, Y = y_j$). The behavior of the classifier can be sampled over N iterated experiments to obtain a count matrix N_{XY} where $N_{XY}(x_i, y_j) = N_{ij}$ counts the number of times that the joint event ($X = x_i, Y = y_j$) occurs. We say that N_{XY} is the (*count-based*) *confusion matrix or contingency table* of the classifier.

Since a confusion matrix is an aggregate recording of the classifier’s decisions, the characterization of the classifier’s performance by means of a measure or set of measures over its confusion matrix is an interesting goal.

One often used measure is *accuracy*, the proportion of times the classifier takes the correct decision $A(N_{XY}) \approx \sum_i N_{XY}(x_i, y_i)/N$. But this has often been deemed biased towards classifiers acting on

non-uniform prior distributions of input patterns (Ben-David, 2007; Sindhwani et al., 2004). For instance, with continuous speech corpora, the *silence* class may account for 40–60% percent of input patterns making a *majority classifier* that always decides $Y = \textit{silence}$, the most prevalent class, quite accurate but useless. Related measures based in proportions over the confusion matrix can be found in Sokolova and Lapalme (2009).

On these grounds, Kononenko and Bratko (1991) have argued for the factoring *out* of the influence of prior class probabilities in similar measures. Yet, Ben-David (2007) has argued for the use of measures that correct naturally for random decisions, like *Cohen’s kappa*, although this particular measure seems to be affected by the marginal distributions.

The *Receiver Operating Characteristic (ROC)* curve (Fawcett, 2006) has often been considered a good *visual* characterization of *binary* confusion matrices built upon proportion measures, but its generalization to higher input and output set cardinals is not as effective. Likewise, an extensive *Area Under the Curve, (AUC)* for a ROC has often been considered an indication of good classifiers (Bradley, 1997; Fawcett, 2006), but the calculation of its higher dimensional analogue, the *Volume Under the Surface, (VUS)* (Hand and Till, 2001) is less manageable. It may also suffer from comparability issues across classifiers (Hand, 2009).

*Corresponding author. Phone: +34 91 624 87 38. Fax: +34 91 624 87 49

Email addresses: fva@tsc.uc3m.es (Francisco J. Valverde-Albacete), carmen@tsc.uc3m.es (Carmen Peláez-Moreno)

Preprint submitted to Pattern Recognition Letters

49 A better ground for discussing performance than
50 count confusion matrices may be empirical estimates of
51 the joint distribution between input and outputs, like the
52 maximum likelihood estimate used throughout this letter
53 $P_{XY}(x_i, y_j) \approx \hat{P}_{XY}^{\text{MLE}}(x_i, y_j) = N(x_i, y_j)/N$. The sub-
54 sequent consideration of the classifier as an analogue
55 of a communication channel between input and output
56 class identifiers enables the importing of information-
57 theoretic tools to characterize the “classification chan-
58 nel”. This technique is already implicit in the work of
59 Miller and Nicely (1955).

60 With this model in mind, Sindhwani et al. (2004) ar-
61 gued for entropic measures that take into account the
62 information transfer through the classifier, like the *ex-*
63 *pected mutual information* between the input and output
64 distributions (Fano, 1961)

$$65 \quad MI_{P_{XY}} = \sum_{x,y} P_{X,Y}(x,y) \log \frac{P_{X,Y}(x,y)}{P_X(x)P_Y(y)} \quad (1)$$

67 and provided a contrived example with three confusion
68 matrices with the same accuracy but clearly differing
69 performances, in their opinion due to differences in mu-
70 tual information. Such examples are alike those put
71 forth by Ben-David (2007) to argue for Cohen’s kappa
72 as an evaluation metric for classifiers.

73 For the related task of clustering, Meila (2007) used
74 the *Variation of Information*, that actually amounts to
75 the sum of their mutually conditioned entropies as a true
76 distance between the two random variables

$$77 \quad VI_{P_{XY}} = H_{P_{X|Y}} + H_{P_{Y|X}}.$$

78
79 In this letter we first try to reach a more complete
80 understanding of what is a good classifier by develop-
81 ing an overall constraint on the total entropy balance
82 attached to its joint distribution. Generalizing over the
83 input and output class set cardinalities will allow us to
84 present a visualization tool in section 2.2 for classifier
85 evaluation that we will further explore in some exam-
86 ples both from real and synthetic data in section 2.3. In
87 section 2.4 we try to extend the tools to unmask major-
88 ity classifiers as bad classifiers. Finally we discuss the
89 affordances of these tools in the context of previously
90 used techniques.

91 2. Information-Theoretic Analysis of Confusion 92 Matrices

93 2.1. The Balance equation and the 2-simplex

94 Let $P_{XY}(x,y)$ be an estimate of the joint proba-
95 bility mass function (pmf) between input and output

96 with marginals $P_X(x) = \sum_{y_j \in Y} P_{X,Y}(x, y_j)$ and $P_Y(y) =$
97 $\sum_{x_i \in X} P_{X,Y}(x_i, y)$.

98 Let $Q_{XY} = P_X \cdot P_Y$ be the pmf¹ with the same
99 marginals as P_{XY} considering them to be independent
100 (that is, describing independent variables). Let $U_{XY} =$
101 $U_X \cdot U_Y$ be the product of the uniform, maximally en-
102 tropic pmfs over X and Y , $U_X(x) = 1/n$ and $U_Y(y) =$
103 $1/p$. Then the loss in uncertainty from U_{XY} to Q_{XY} is
104 the difference in entropies:

$$105 \quad \Delta H_{P_X \cdot P_Y} = H_{U_X \cdot U_Y} - H_{P_X \cdot P_Y} \quad (2)$$

107 Intuitively, $\Delta H_{P_X \cdot P_Y}$ measures how far the classifier
108 is operating from the most general situation possible
109 where all inputs are equally probable, which prevents
110 the classifier from specializing in an overrepresented
111 class to the detriment of classification accuracy in oth-
112 ers. Since $H_{U_X} = \log n$ and $H_{U_Y} = \log p$, $\Delta H_{P_X \cdot P_Y}$ may
113 vary from $\Delta H_{P_X \cdot P_Y}^{\min} = 0$, when the marginals themselves
114 are uniform $P_X = U_X$ and $P_Y = U_Y$, to a maximum
115 value $\Delta H_{P_X \cdot P_Y}^{\max} = \log n + \log p$, when they are Kronecker
116 delta distributions.

117 We would like to relate this entropy decrement to the
118 expected mutual information $MI_{P_{XY}}$ of a joint distribu-
119 tion. For that purpose, we realize that the mutual in-
120 formation formula (1) describes the decrease in entropy
121 when passing from distribution $Q_{XY} = P_X \cdot P_Y$ to P_{XY}

$$122 \quad MI_{P_{XY}} = H_{P_X \cdot P_Y} - H_{P_{XY}}. \quad (3)$$

124 And finally we invoke the well-known formula relating
125 the joint entropy $H_{P_{XY}}$ and the expected mutual infor-
126 mation $MI_{P_{XY}}$ to the conditional entropies of X given Y ,
127 $H_{P_{X|Y}}$ (Y given X , $H_{P_{Y|X}}$ respectively):

$$128 \quad H_{P_{XY}} = H_{P_{X|Y}} + H_{P_{Y|X}} + MI_{P_{XY}} \quad (4)$$

130 Therefore $MI_{P_{XY}}$ may range from $MI_{P_{XY}}^{\min} = 0$ when
131 $P_{XY} = P_X \cdot P_Y$, a bad classifier, to a theoretical max-
132 imum $MI_{P_{XY}}^{\max} = (\log n + \log p)/2$ in the case where the
133 marginals are uniform and input and output are com-
134 pletely dependent, an excellent classifier.

135 Recall the variation of information definition in Eq.
136 (5).

$$137 \quad VI_{P_{XY}} = H_{P_{X|Y}} + H_{P_{Y|X}} \quad (5)$$

139 For optimal classifiers with deterministic relation from
140 the input to the output, and diagonal confusion matrices
141 $VI_{P_{XY}}^{\min} = 0$, e.g., all the information about X is borne by

¹We drop the explicit variable notation in the distributions from
now on.

142 Y and vice versa. On the contrary, when they are inde- 189
 143 pendent $VI_{P_{XY}}^{\max} = H_{P_X} + H_{P_Y}$, the case with inaccurate 190
 144 classifiers which uniformly redistribute inputs among 191
 145 all outputs. 192

146 Collecting Eqs. (2)–(5) results in the *balance equa-* 193
 147 *tion for information related to a joint distribution*, our 194
 148 first result, 195

$$149 \quad H_{U_{XY}} = \Delta H_{P_X \cdot P_Y} + 2MI_{P_{XY}} + VI_{P_{XY}} \quad (6) \quad 196$$

151 The balance equation suggests an *information dia-* 197
 152 *gram* somewhat more complete than what is normally 198
 153 used for the relations between the entropies of two vari- 199
 154 ables as depicted in Fig. 1(a) (compare to Yeung, 1991, 200
 155 Fig. 1). In this diagram we distinguish the familiar de- 201
 156 composition of the joint entropy $H_{P_{XY}}$ as the two en- 202
 157 tropies H_{P_X} and H_{P_Y} whose intersection is $MI_{P_{XY}}$. But 203
 158 notice that the increment between $H_{P_{XY}}$ and $H_{P_X \cdot P_Y}$ is 204
 159 yet again $MI_{P_{XY}}$, hence the expected mutual informa- 205
 160 tion appears *twice* in the diagram. Further, the interior 206
 161 of the outer rectangle represents $H_{U_X \cdot U_Y}$, the interior of 207
 162 the inner rectangle $H_{P_X \cdot P_Y}$ and $\Delta H_{P_X \cdot P_Y}$ represents their 208
 163 difference in areas. The absence of the encompassing 209
 164 outer rectangle in Fig. 1a was specifically puzzled at by 210
 165 Yeung (1991). 211

166 Notice that, since both U_X and U_Y on the one hand 212
 167 and P_X and P_Y are independent as marginals of U_{XY} and 213
 168 Q_{XY} , respectively, we may write: 214

$$169 \quad \Delta H_{P_X \cdot P_Y} = (H_{U_X} - H_{P_X}) + (H_{U_Y} - H_{P_Y}) = \Delta H_{P_X} + \Delta H_{P_Y} \quad 215$$

$$170 \quad (7) \quad 216$$

171 where

$$172 \quad \Delta H_{P_X} = H_{U_X} - H_{P_X} \quad \Delta H_{P_Y} = H_{U_Y} - H_{P_Y} \quad (8) \quad 217$$

174 This and the occurrence of twice the expected mutual 220
 175 information in Eq. (6) suggests a different information 221
 176 diagram, depicted in Fig. 1(b). Both variables X and 222
 177 Y now appear somehow decoupled—in the sense that 223
 178 the areas representing them are disjoint—yet there is a 224
 179 strong coupling in that the expected mutual information 225
 180 appears in both H_{P_X} and H_{P_Y} . This suggests writing 226
 181 separate balance equations for each variable, to be used 227
 182 in Sec. 2.4, 228

$$183 \quad H_{U_X} = \Delta H_{P_X} + MI_{P_{XY}} + H_{P_{X|Y}} \quad H_{U_Y} = \Delta H_{P_Y} + MI_{P_{XY}} + H_{P_{Y|X}} \quad 229$$

$$184 \quad (9) \quad 230$$

185 Our interpretation for the balance equation is that the 231
 186 “raw” uncertainty available in U_{XY} minus the deviation 232
 187 of the input data from the uniform distribution ΔH_{P_X} , a 233
 188 given, is redistributed in the classifier-building process 234
 235

to the information being transferred from input to out- 190
 put $MI_{P_{XY}}$. This requires as much mutual information to 191
 stochastically bind the input to the output, thereby trans- 192
 forming $P_X \dot{P}_Y$ into P_{XY} , and incurs in an uncertainty 193
 decrease at the output equal to ΔH_{P_Y} . The residual un- 194
 certainty $H_{P_{X|Y}} + H_{P_{Y|X}}$ should measure how efficient the 195
 process is: the smaller, the better. 196

To gain further understanding of the entropy decom- 197
 position suggested by the balance equation, from Eq. 198
 (6) and the paragraphs following Eqs. (2)–(5), we obtain 199

$$200 \quad H_{U_{XY}} = \Delta H_{P_X \cdot P_Y} + 2MI_{P_{XY}} + VI_{P_{XY}} \quad 201$$

$$202 \quad 0 \leq \Delta H_{P_X \cdot P_Y}, 2MI_{P_{XY}}, VI_{P_{XY}} \leq H_{U_{XY}} \quad 203$$

imposing severe constraints on the values the quantities 204
 may take, the most conspicuous of which is that given two 205
 of the quantities the third one is fixed. Normalizing 206
 by $H_{U_{XY}}$ we get

$$207 \quad 1 = \Delta H'_{P_X \cdot P_Y} + 2MI'_{P_{XY}} + VI'_{P_{XY}} \quad (10)$$

$$208 \quad 0 \leq \Delta H'_{P_X \cdot P_Y}, 2MI'_{P_{XY}}, VI'_{P_{XY}} \leq 1. \quad 209$$

210 This is the 2-simplex in normalized $\Delta H'_{P_X \cdot P_Y} \times$
 211 $2MI'_{P_{XY}} \times VI'_{P_{XY}}$ space depicted in Fig. 2(a), a three-
 212 dimensional representation of classifier performance: 213
 each classifier with joint distribution P_{XY} can be char- 214
 215 acterized by its *joint entropy fractions*, $F_{XY}(P_{XY}) =$
 $[\Delta H'_{P_{XY}}, 2 \times MI'_{P_{XY}}, VI'_{P_{XY}}]$.

2.2. De Finetti entropy diagrams

216 Since the ROC curve is a bi-dimensional character- 217
 218 ization of binary confusion matrices we might wonder 219
 if the constrained plane above has a simpler visualiza- 220
 tion. Consider the 2-simplex in Eq. (10) and Fig. 2(a). 221
 Its projection onto the plane with director vector is 222
 $(1, 1, 1)$ is its *de Finetti (entropy) diagram*, represented 223
 in Fig. 2(b). Alternatively to the three-dimensional rep- 224
 225 resentation, each classifier can be represented as a point 226
 at coordinates F_{XY} in the de Finetti diagram. 227

The de Finetti entropy diagram shows as an equilat- 228
 eral triangle, hence the alternative name *entropy trian-* 229
 230 *gle*, each of whose sides and vertices represents classi- 231
 232 fier performance-related *qualities*: 233

- If P_X and P_Y are independent in $Q_{XY} = P_X \cdot P_Y$ then 234
 $F_{XY}(Q_{XY}) = [\cdot, 0, \cdot]$. The lower side is the geomet- 235
 236 ric locus of distributions with no mutual informa- 237
 238 tion transfer between input and output: the closer 239
 239 a classifier is to this side, the more unreliable the 240
 240 classifier decisions are. 241




Figure 1: **Extended information diagrams of entropies related to a bivariate distribution:** the expected mutual information appears *twice*. (a) Extended diagram, and (b) Modified extended diagram.

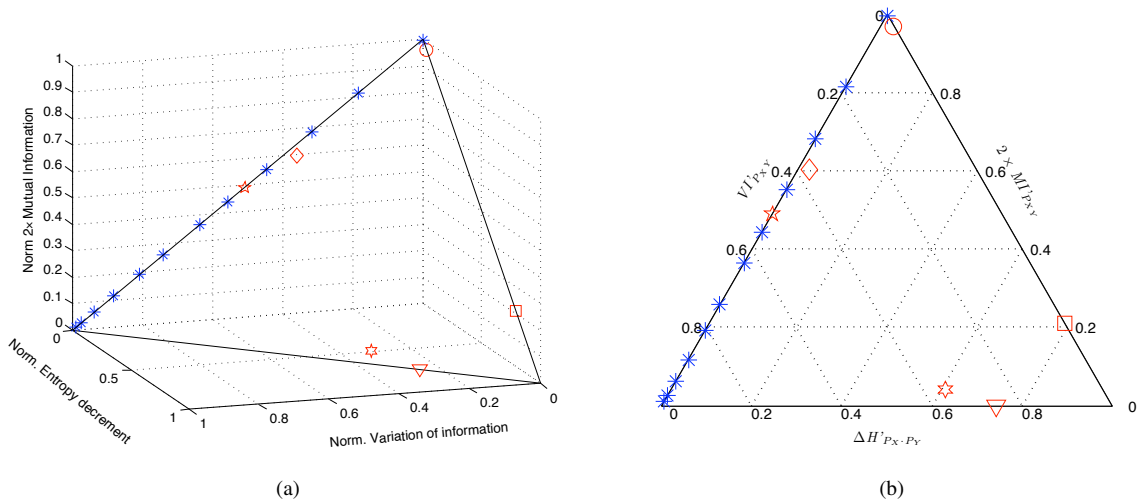


Figure 2: (color on-line) **Entropic representations for bivariate distribution of the synthetic examples of Fig. 3:** (a) The 2-simplex in three-dimensional, normalized entropy space $\Delta H'_{P_X \cdot P_Y} \times VI'_{P_{XY}} \times 2MI'_{P_{XY}}$ and (b) the de Finetti entropy diagram or entropy triangle, a projection of the 2-simplex onto a two-dimensional space (explanations in Sec. 2.3).

- If the marginals of P_{XY} are uniform $P_X = U_X$ and $P_Y = U_Y$ then $F_{XY}(P_{XY}) = [0, \cdot, \cdot]$. This is the locus of classifiers that are not trained with over-represented classes and therefore cannot specialize in any of them: the closer to this side, the more generic the classifier.
- Finally, if P_{XY} is a diagonal matrix, then $P_X = P_Y$ and $F_{XY}(P_{XY}) = [\cdot, \cdot, 0]$. The right-hand side is the region of classifiers with no variation of information, that is, no remanent information in the conditional entropies: this characterizes classifiers which transfer as much information from H_{P_X} to

H_{P_Y} as they can.

Moving away from these sides the corresponding magnitudes grow until saturating at the opposite vertices, which therefore represent ideal, *prototypical classifier loci*:

- The upper vertex $F_{XY}(optimal) = [0, 1, 0]$ represents *optimal classifiers* with the highest information transfer from input to output and highly entropic priors.
- The vertex to the left $F_{XY}(inaccurate) = [0, 0, 1]$ represents *inaccurate classifiers*, with low

information-transfer with highly entropic priors.

- The vertex to the right $F_{XY}(\text{underperforming}) = [0, 0, 1]$ represents *underperforming classifiers*, with low information transfer and low-entropic priors either at an easy task or refusing to deliver performance.

In the next section we develop intuitions over the de Finetti diagram by observing how typical examples, real and synthetic appear in it.

But first we would like to extend it theoretically to cope with the separate information balances of the marginal distributions. Recall that the modified information diagram in Fig. 1(b) suggest a decoupling of the information flow from input to output further supported by eq. 9. These describe the *marginal fractions* of entropy when the normalization is done with H_{U_X} and H_{U_Y} respectively

$$\begin{aligned} F_X(P_{XY}) &= [\Delta H'_{P_X}, MI'_{P_{XY}}, VI'_X = H'_{P_{XY}}] \\ F_Y(P_{XY}) &= [\Delta H'_{P_Y}, MI'_{P_{XY}}, VI'_Y = H'_{P_{Y|X}}] \end{aligned} \quad (11)$$

hence we may consider the de Finetti marginal entropy diagrams for both F_X and F_Y to visualize the entropy changes from input to output.

Furthermore, since the normalization factors involved are directly related to those in the joint entropy balance, and the $MI'_{P_{XY}}$ has the same value in both marginal diagrams when $n = p$, we may represent the fractions for F_X and F_Y side by side those of F_{XY} in an *extended de Finetti entropy diagram*: the point F_{XY} , being an average of F_X and F_Y , will appear in the diagram flanked by the latter two. We show in Sec. 2.4 examples of such extended diagrams and their use.

2.3. Examples

To clarify the usefulness of our tools in assessing classifier performance we explored data from real classifiers and synthetic examples to highlight special behaviors.

First, consider:

- matrices a , b , and c from Sindhvani et al. (2004), reproduced with the same name in Fig. 3,
- a matrix whose marginals are closer to a uniform distribution, a matrix whose marginals are closer to a Kronecker delta, and the confusion matrix of a majority classifier, with a delta output distribution but a more spread input distribution—matrices d , e and f in Fig. 3 respectively—, and

- a series of distributions obtained by convex combination $P_{XY} = (1 - \lambda) \cdot (P_X \cdot P_Y) + \lambda \cdot (P_{X=Y})$ from a uniform bivariate $(P_X \cdot P_Y)$ to a uniform diagonal $(P_{X=Y})$ distribution as the combination coefficient λ ranges in $[0, 1]$.

The contrived examples in Sindhvani et al. (2004), matrices a , b , and c —represented in both diagrams in Fig. 2 as a diamond, a pentagram, and a hexagram, respectively—pointed out there at a need for new performance metrics, since they all showed the same accuracy. The diagrams support the intuition that matrix a describes a slightly better classifier than matrix b which describes a better classifier than matrix c (see Sec. 2.4 for a further analysis of the behavior of c).

Figs. 2(a) and 2(b) demonstrate that there are clear differences in performance between a classifier with more uniform marginals and one with marginals more alike Kronecker deltas (matrices d and e in Fig. 3, the circle and square, respectively). Furthermore, an example of a majority classifier (matrix f , the downwards triangle) shows in the diagram as underperforming: it will be further analyzed in Sec. 2.4.

From the convex combination we plotted the line of asterisks at $\Delta H'_{P_X \cdot P_Y} = 0$ in Figs. 2(a) and (b). When the interpolation coefficient for the diagonal is null, we obtain the point at $\Delta H'_{P_X \cdot P_Y} = 0, VI'_{P_X \cdot P_Y} = 0$ for the worst classifier. As the coefficient increases, the asterisks denote better and better hypothetical classifiers until reaching the apex of the triangle, the best. We simulated in this guise the estimation of classifiers in improving SNR ratios for each point in the line, as shown below on real data.

In order to appraise the usefulness of the representation on real data we visualized in Fig. 4(a) the performance of several series of classifiers. The circles to the right describe a classical example of the performance of human listeners in a 16-consonant human-speech recognition task at different SNR (Miller and Nicely, 1955). They evidence the outstanding recognition capabilities of humans, always close to maximum available information transfer at $VI'_{P_{XY}} = 0$, with a graceful degradation as the available information decreases with decreasing SNR—from 12dB at the top of the line to -18dB at the bottom. And they also testify to the punctiliousness of those authors' in keeping to maximally generic input and output distributions at $\Delta H'_{P_X \cdot P_Y} \approx 0$.

On the other hand, the asterisks, plus signs and crosses are series of automatic speech recognizers using the SpeechDat database (Moreno, 1997). They motivated this work in characterizing classifiers by means of entropic measures.

$$a = \begin{bmatrix} 15 & 0 & 5 \\ 0 & 15 & 5 \\ 0 & 0 & 20 \end{bmatrix} \quad b = \begin{bmatrix} 16 & 2 & 2 \\ 2 & 16 & 2 \\ 1 & 1 & 18 \end{bmatrix} \quad c = \begin{bmatrix} 1 & 0 & 4 \\ 0 & 1 & 4 \\ 1 & 1 & 48 \end{bmatrix}$$

$$d = \begin{bmatrix} 15 & 0 & 0 \\ 0 & 18 & 0 \\ 0 & 0 & 27 \end{bmatrix} \quad e = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 57 \end{bmatrix} \quad f = \begin{bmatrix} 0 & 0 & 5 \\ 0 & 0 & 5 \\ 0 & 0 & 50 \end{bmatrix}$$

Figure 3: **Examples of synthetic confusion matrices with varied behavior:** a , b and c from (Sindhwani et al., 2004), d a matrix whose marginals tend towards uniformity, e a matrix whose marginals tend to Kronecker’s delta and f the confusion matrix of a majority classifier.




Figure 4: (color on-line) **Examples of use of the de Finetti entropy diagram to assess classifiers:** (a) human and machine classifier performance in consonant recognition tasks, and (b) the performance of some prototypical communication channel models.

- 356 • The series of squares describes a 18-class phonetic 374
357 recognition task with worsening SNR that does not 375
358 use any lexical information. This is roughly compar- 376
359 able to the experiments in Miller and Nicely 377
360 (1955) and highlights the wide gap at present be- 378
361 tween human and machine performance in pho- 379
362 netic recognition.
- 363 • The series of plus signs describes phonetic confu- 380
364 sions on the same phonetic recognition task when 381
365 lexical information is incorporated. Notice that 382
366 the tendency in either series is not towards the 383
367 apex of the entropy triangle, but towards increas- 384
368 ing $\Delta H'_{P_X P_Y}$, suggesting that the learning technique 385
369 used to build the classifiers is not making a good 386
370 job of extracting all the phonetic information avail- 387
371 able from the data, choosing to specialize the clas- 388
372 sifier instead. Further, the additional lexical con- 389
373 straints on their own do not seem to be able to span 390
391 the gap with human performance.
- 392 • Finally, the asterisks describe a series of classifiers 393
394 for a 10-digit recognition task on the same data. 395
396 The very high values of all the coordinates suggest 397
398 that this is a well-solved task at all those noise con- 399
400 ditions.

Notice that, although all these tasks have different class set cardinalities, they can be equally well compared in the same entropy triangle.

Since the simplex was developed for joint distributions, other objects characterized by these, such as *communication channel models*, may also be explored with the technique. These are high level descriptions of the end-to-end input and output-symbol classification capabilities of a communication system. Fig. 4(b) depicts three types of channels from MacKay (2003):

- the *binary symmetric channel* with $n = p = 2$ where we have made the probability of error range

in $p_e \in [0, 0.5]$ in 0.05 steps to obtain the series plotted with asterisks,

- the *binary erasure channel* with $n = 2$; $p = 3$ with the erasure probability ranging in $p_e \in [0, 1.0]$ in 0.01 steps plotted with circles, and
- the *noisy typewriter* with $n = p = 27$ describing a typewriter with a convention on the errors it commits, plotted as a pentagram.

As channels are actually defined by conditional distributions $P_{Y|X}(y|x)$ we multiplied them with a uniform prior $P_X = U_X$ to plot them. Although $P_X = U_X$ the same cannot be said of P_Y what accounts for the fact that on most of the sample points in the binary erasure channel we have $\Delta H'_{P_X \cdot P_Y} \neq 0$. On the other hand, the symmetries in the binary symmetric channel and the noisy typewriter account for $\Delta H'_{P_X \cdot P_Y} = 0$.

Notice how in the entropy triangle we can even make sense of a communication channel with different input and output symbol set cardinalities, e.g. the binary erasure channel.

2.4. De Finetti diagram analysis of majority classifiers

Majority classifiers are capable of achieving a very high accuracy rate but are of limited interest. It is often required that good performance evaluation measures for classifiers show a baseline both for random and majority classifiers (Ben-David, 2007). For instance, majority classifiers should:

- have a low output entropy, a high $\Delta H'_{P_Y}$, whatever its $\Delta H'_{P_X}$ value.
- have a low information transfer $MI'_{P_{XY}}$.
- have some output conditional entropy, hence some $VI'_{P_{XY}}$.

Matrix f in Fig. 3 is the confusion matrix of majority classifier with a non-uniform input marginal. We would like to know whether this behavior could be gleaned from a de Finetti diagram.

In Fig. 5(a) we have plotted again the joint entropy fractions for the synthetic cases analyzed above, together with the entropy fractions of their marginals. For most of the cases, all three points coincide—showing as a crosses within circles.

But matrices a , c and f —diamond, hexagram and downwards triangle in Fig. 5(a)—show differences in joint and marginal fractions. The most striking cases are those of matrices c and f , whose uncertainty diminishes dramatically from input to output.

Matrix f in Fig. 3 models the behavior of a majority classifier with the same input marginal as a – e . The marginal fraction points appear flanking this, at $F_X(f) = [0.45, 0, 0.55]$ and $F_Y(f) = [1, 0, 0]$. The accuracy for this classifier would be around 0.83.

In a sense, this classifier is cheating: without any knowledge of the actual classification instances it has optimized the average accuracy, but will be defeated if the input distribution gets biased towards a different class in the deployment (test) phase. It is now quite clear that c , being close to a majority classifier, attains its accuracy by specialization too.

Indeed, observing matrix a we may pinpoint the fact that its zero-pattern seems to be the interpolation of a diagonal confusion matrix and the confusion matrix of a majority classifier. This fact shows as the two flanking marginal fractions to the diamond at approximately $F_{XY}(a) = [0.03, 0.6, 0.37]$ in Fig. 5. However, since P_X was wisely kept uniform, $\Delta H'_{P_X} = 0$ at $F_Y(a) = [0, 0.6, 0.4]$ the classifier could only specialize to $F_Y(a) = [0.06, 0.6, 0.34]$.

These examples suggest that:

- *Specialization is a reduction in $VI'_{P_{XY}}$ caused by the reduction in VI'_{P_Y} brought about by the increase in $\Delta H'_{P_Y}$, that is manipulation of the output marginal distribution.*
- *Classifiers with diagonal matrices $VI'_{P_{XY}} = 0$ need not (and classifiers with uniform marginals $\Delta H'_{P_{XY}} = 0$ cannot) specialize.*
- *Maintaining uniform input marginals amounts to a sort of regularization preventing specialization further from transforming all $\Delta H'_{P_Y}$ into a decrement of $VI'_{P_{XY}}$.*

For real classifiers, we have plotted in Fig. 5(b) the marginal fractions of all the classifiers in Fig. 4(a). Again, for most of them, the marginal fractions coincide with the joint fractions. But for the phonetic SpeechDat task plotted with squares we observe how with decreasing SNR the classifier has to resort to specialization. With increasing SNR it can concentrate on increasing the expected mutual information transmitted from input to output.

3. Discussion and conclusions

We have provided a mathematical tool to analyze the behavior of multi-class classifiers by means of the balance of entropies of the joint probability mass distribution of input and output classes as estimated from their confusion matrix or contingency table.




Figure 5: (color on-line) **Extended de Finetti entropy diagrams for synthetic and real examples:** (a) for the synthetic confusion matrices of Fig. 2(b), and (b) for the real confusion matrices of Fig. 4(a). The expected mutual information coordinate is maintained in the three points for each confusion matrix.

486 The balance equation takes into consideration the 516
 487 Kullback-Leibler divergence between the uniform and 517
 488 independent distributions with the same marginals as 518
 489 the original one, twice the expected mutual infor- 519
 490 mation between the independent and joint distribu- 520
 491 tions with identical marginals—also a Kullback-Leibler 521
 492 divergence—and the variation of information, the differ- 522
 493 ence between the joint entropy and the expected mutual 523
 494 information. 524

495 This balance equation can either be visualized as 525
 496 the 2-simplex in three-dimensional entropy space, with 526
 497 dimensions being normalized instances of those men- 527
 498 tioned above; or it can be projected to obtain a ternary 528
 499 plot, a conceptual diagram for classifiers resembling a 529
 500 triangle whose vertices characterize optimal, inaccurate, 530
 501 or underperforming classifiers. 531

502 Motivated by the need to explain the accuracy- 532
 503 improving behavior of majority classifiers we also intro- 533
 504 duced the extended de Finetti entropy diagram where 534
 505 input and output marginal entropy fractions are visual- 535
 506 ized side by side the joint entropy fractions. This al- 536
 507 lows us to detect those classifiers resorting to special- 537
 508 ization to increase their accuracy without increasing the 538
 509 mutual information. It also shows how this behavior 539
 510 can be limited by maintaining adequately uniform input 540
 511 marginals. 541

512 We have used these tools to visualize confusion ma- 542
 513 trices for both human and machine performance in sev- 543
 514 eral tasks of different complexities. The balance equa- 544
 515 tion and de Finetti diagrams highlight the following 545

516 facts:

- The expected mutual information transmitted from input to output is limited by the need to use as much entropy to bind together in stochastic dependency both variables $MI_{P_{XY}} \leq H_{P_X, P_Y}/2$.
- Even when the mutual information between input and output is low, if the marginals have in-between uncertainty $0 < \Delta H_{P_{XY}} < \log n + \log p$ and $P_X \neq P_Y$, a classifier may become *specific*—e.g. specialize in overrepresented classes—to decrease the variation of information, effectively increasing its accuracy.
- The variation of information is actually the information *not* being transmitted by the classifier, that is, the uncoupled information between input and output. This is a good target for improving accuracy without decreasing the *genericity* of the resulting classifier, e.g., its non-specificity.

534 All in all the three leading assertions contextualize 535
 535 and nuance the assertion in Sindhvani et al. (2004), viz. 536
 536 the higher the mutual information, the more generic and 537
 537 accurate (less specialized and inaccurate) a classifier’s 538
 538 performance will be.

539 The generality and applicability of the techniques 540
 540 have been improved by using information-theoretic 541
 541 measures that pertain not only to the study of confu- 542
 542 sion matrices but, in general, to bivariate distributions 543

543 such as communication channel models. However, the
544 influence of the probability estimation method is as yet
545 unexplored. Unlike Meila (2007), we have not had to
546 suppose equality of sets of events in the input or output
547 spaces or symmetric confusion matrices.

548 Comparing the de Finetti entropy diagram and the
549 ROC is, at best, risky for the time being. On the one
550 hand, the ROC is a well-established technique that af-
551 fords a number of intuitions in which practitioners are
552 well-versed, including a rather direct relation to accu-
553 racy. Also, the VUS shows promise of actually be-
554 coming a figure-of-merit for multi-class classifiers. For
555 a more widespread use, the entropy triangle should
556 offer such succinct, intuitive affordances too. In the
557 case of accuracy, we intend to use Fano’s inequality
558 to bridge our understanding of proportion- and entropy-
559 based measures.

560 On the other hand, the ROC only takes into consid-
561 eration those judgments of the classifier *within* the joint
562 entropy area in the Information Diagram and is thus un-
563 able to judge how close to genericity is the classifier,
564 unlike the $\Delta H'_{P_X \cdot P_Y}$ coordinate of the entropy triangle.
565 Likewise, the ROC has so far been unable to obtain the
566 result that as much information as actually transmitted
567 from input to output must go into creating the stochastic
568 dependency between them.

569 To conclude, however suggestive aggregate measures
570 like entropy or mutual information may be for captur-
571 ing at a glance the behavior of classifiers, they offer lit-
572 tle in the way of analyzing the actual classification er-
573 rors populating their confusion matrices. We believe the
574 analysis of mutual information as a random variable of a
575 bivariate distribution (Fano, 1961, pp. 27–31) may offer
576 more opportunities for *improving* classifiers as opposed
577 to *assessing* them.

578 Acknowledgements

579 This work has been partially supported by the Span-
580 ish Government-Comisión Interministerial de Cien-
581 cia y Tecnología projects 2008-06382/TEC and 2008-
582 02473/TEC and the regional projects S-505/TIC/0223
583 (DGUI-CM) and CCG08-UC3M/TIC-4457 (Comu-
584 nidad Autónoma de Madrid - UC3M).

585 The authors would like to thank C. Bousoño-Calzón
586 and A. Navia-Vázquez for comments on early versions
587 of this paper and A. I. García-Moral for providing the
588 confusion matrices from the automatic speech recogniz-
589 ers.

590 References

- 591 Ben-David, A., 2007. A lot of randomness is hiding in accuracy. En-
592 gineering Applications of Artificial Intelligence 20 (7), 875–885.
- 593 Bradley, A. P., 1997. The use of the area under the ROC curve in
594 the evaluation of machine learning algorithms. Pattern Recognition
595 30 (7), 1145–1159.
- 596 Fano, R. M., 1961. Transmission of Information: A Statistical Theory
597 of Communication. The MIT Press.
- 598 Fawcett, T., 2006. An introduction to ROC analysis. Pattern Recogni-
599 tion Letters 27 (8), 861–874.
- 600 Hand, D. J., 2009. Measuring classifier performance: a coherent alter-
601 native to the area under the ROC curve. Machine Learning 77 (1),
602 103–123.
- 603 Hand, D. J., Till, R. J., 2001. A simple generalisation of the Area
604 Under the ROC Curve for multiple class classification problems.
605 Machine Learning 45, 171–186.
- 606 Kononenko, I., Bratko, I., 1991. Information-based evaluation crite-
607 rion for classifier’s performance. Machine Learning 6, 67–80.
- 608 MacKay, D. J., 2003. Information Theory, Inference, and Learning
609 Algorithms. Cambridge University Press.
- 610 Meila, M., 2007. Comparing clusterings—an information based dis-
611 tance. Journal of Multivariate Analysis 28, 875–893.
- 612 Miller, G. A., Nicely, P. E., 1955. An analysis of perceptual confu-
613 sions among some english consonants. The Journal of the Acoustic
614 Society of America 27 (2), 338–352.
- 615 Moreno, A., 1997. SpeechDat Spanish Database for Fixed Telephone
616 Network. Tech. rep., Technical University of Catalonia, Barcelona,
617 Spain.
- 618 Sindhwani, V., Rakshit, S., Deodhare, D., Erdogmus, D., Principe, J.,
619 Niyogi, P., 2004. Feature selection in MLPs and SVMs based on
620 maximum output information. IEEE Transactions on Neural Net-
621 works 15 (4), 937–948.
- 622 Sokolova, M., Lapalme, G., Jul 2009. A systematic analysis of perfor-
623 mance measures for classification tasks. Information Processing &
624 Management 45 (4), 427–437.
- 625 Yeung, R., 1991. A new outlook on Shannon’s information measures.
626 IEEE Transactions on Information Theory 37 (3), 466–474.