INDEPENDENT COMPONENTS OF EXPERTISE

Heikki Hyötyniemi and *Vesa Hasu*

Helsinki University of Technology, Control Engineering Laboratory,
Konemiehentie 2, FIN-02150 Espoo, Finland

This paper presents some results that have been found when a new view of human expertise is employed: it is assumed that the "chunks" of expert knowledge are based on "independent components" in the high-dimensional observation space. These components are exploited by the optimised expert model because of their power of representing the data; and because of their relevance they can also be detected and extracted using automatic statistical tools. As an application example, modelling of the visual outlook of the so called flotation process are presented: it seems that the extracted froth prototypes match the sub-symbolic categorization of the human process operator.

1. INTRODUCTION

The traditional view of human expertise assumes that human knowledge is based on symbolic constructs: the inference relies on rules, and representation on conceptual categories or chunks [10,11]. Even though this view has lately been challenged (for that new views see [3]), the basic paradoxes still remain: for example, how can the shift from novice to expert be explained? If expertise were characterized by larger and larger sets of rules or other constructs being gathered, how could the reasoning become faster and more automated along learning?

It turns out that when the expertise is seen as an optimised model of the domain field (in the systems theoretical sense), the traditional problems can be efficiently attacked. For example, in [5], this modelling view is applied to the chunking in chess: it turns out that the "numerical chunks" can essentially relieve the expert memory requirements. The model structure is assumed to be based on *clusters* of data, and the fine structure within the clusters consists of *independent components* (see [4] and [6]). In this paper, this approach is applied to a rather practical field of expert knowledge: analysis of process state in an industrial concentrator plant.

2. INDEPENDENT COMPONENTS

It has been recognized that *independent components* offer a promising way of finding the underlying structure behind observations (see [2]). Independent components are latent variables that span the space of measurements in a physically "meaningful" way; they are subspace axes that make it possible to compress the observation data in a plausible way. Independent component analysis (ICA) is a rather new field of research – the traditional methods like principal component analysis (PCA) are mathematically more efficient, but physically less plausible: they are usually based on *orthogonality* of the latent structure.

There exist specialized algorithms for carrying out independent component analysis; traditionally, the non-Gaussianity of the data distribution is used as the criterion. However, in practical applications, this approach suffers from the non-ideal measurement conditions in the process: the "outliers" that are typical to process measurements may ruin the data analysis. What is more, this approach does not support complex data, because it cannot easily be combined with clustering: non-Gaussianity assumption suggests distorted clusters, whereas, on the other hand, clustering creates new boundaries within the data that may twist the extracted independent components. In this experiment a special algorithm called GGHA [7] is used, where the search of independent components is based on the explicit assumption of *sparsity* of the fine structure, and redundancy between the clusters in the data [14].

It has been shown that the visual cortex applies independent component analysis to visual images, first extracting features like border lines, etc., from the images [15]; in the auditory cortex, analogous results have been detected. As is shown in this paper, the same principles might apply on the higher levels of mental processing.

3. EXAMPLE: FLOTATION PROCESS

As a practical application environment, modelling of an industrial *flotation process* [12] is carried out. In flotation, minerals are purified by letting air bubbles penetrate through the ore pulp: the mineral grains stick to the bubbles, and they are collected on the froth layer on the surface where the concentrate can be collected. The complexity of the flotation process is partly caused by the lack of appropriate measurement devices, and it is still the human operator that is needed to analyse the process state using visual inspection.

The process data was collected at Outokumpu Oy zinc concentrator plant in Pyhäsalmi, Finland. The process has been monitored over several months, resulting in large amounts of measurement data consisting of process variables as well as variables characterizing the visual outlook of the froth layer. The froth features include the intensity distributions in the reflected light, morphological features (bubble size and shape distributions), and measures for the dynamic properties of the froth (bubble collapse rate, froth speed, etc.). The 12-dimensional data vectors are now constructed as follows:

```
1:   maximum peak of the cross-correlation matrix between an image pair
2:   speed of the froth
3:   mean size of bubbles
4:   mean eccentricity (how elliptic the bubbles are)
5:   standard deviation of bubble size
6:   standard deviation of eccentricity
7:   mean red value
8:   standard deviation of red value
9:   skewness of red value (third cumulant of the distribution)
10:  kurtosis of red value (fourth cumulant of the distribution)
11:  mean hue value
12:  standard deviation of hue value
```

Various research reports exist where machine vision has been used in the analysis and classification of flotation froths (for example, see [1] and [13]). However, in these experiments the froth classes are often rather trivial (the classes being something like "big bubbles"/"small bubbles" – see Figs. 1 and 2), and they do not necessarily match the experts' views of what are the relevant phenomena. Now, on the other hand, the bubble size and shape information is not used directly for determining the froth classes; these measures are highly redundant, and the underlying "latent" relationships searched instead. It must be recognized that process properties vary from plant to plant – there does not exist any "universally applicable" classification: only the experienced operators know what is *relevant* in each case.

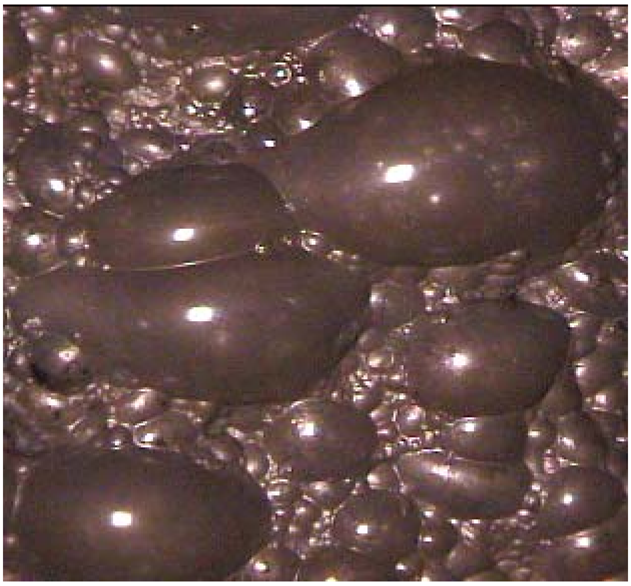

Figure 1. “Big bubbles”

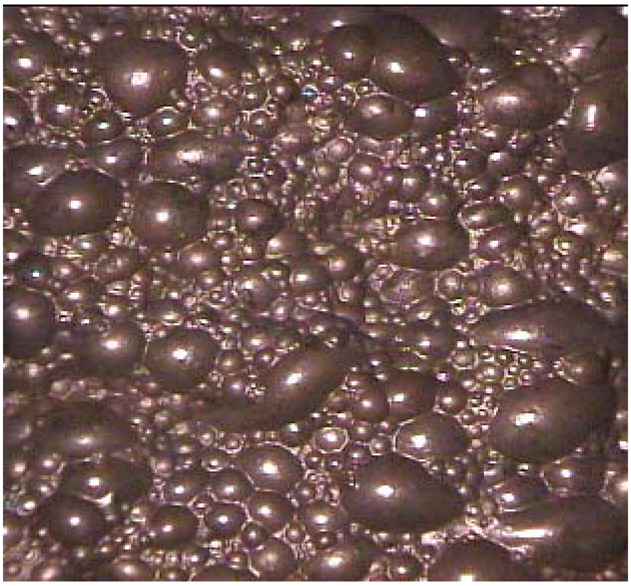

Figure 2. “Small bubbles”

At the Pyhäsalmi plant, an operator enquiry was carried out [9]. Even though it is a well-known problem that expert knowledge cannot be easily extracted, now it was found out that the operators felt that three different froth appearances were particularly important: these main classes were “stiff”, “dry”, and “wet” froth (see Figs. 3, 4, and 5). According to the operator interviews, a special analyser was constructed for automatically extracting the level of “stiffness”, etc., for different images. For example, there is high probability for stiff froth if the froth speed is low and bubble size is small. Although wet and dry sound like being opposite concepts, they are not; there is correlation (that is, they are not orthogonal) but still they have been detected as being independent.

A large set of images were modeled using statistical tools – GGHA analysis, and for reference, also PCA. First, the number of three clusters were chosen, and the fine structure analysis was thereafter carried out. The extracted models were tested using another set of images, and the results are shown in Figs. 6, 7, and 8, and the results are analyzed in Table 1.

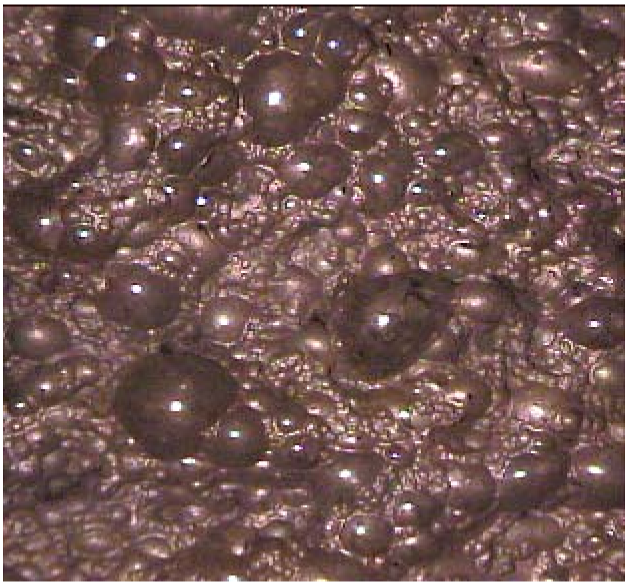

Figure 3. "Stiff" froth

Figure 4. "Dry" froth

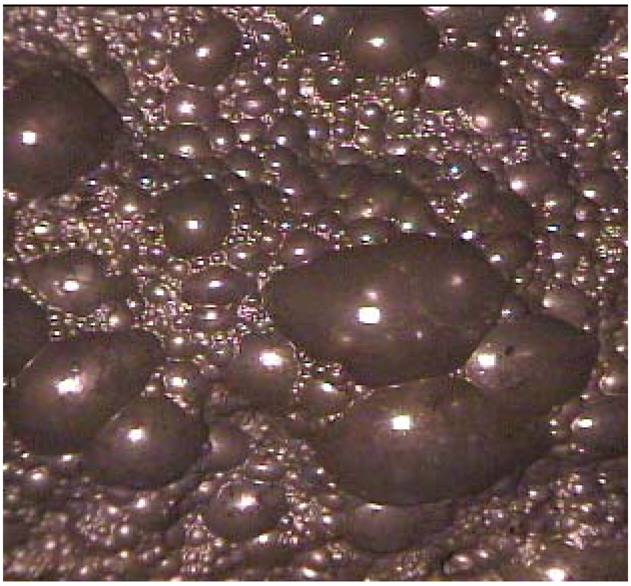

Figure 5. "Wet" froth

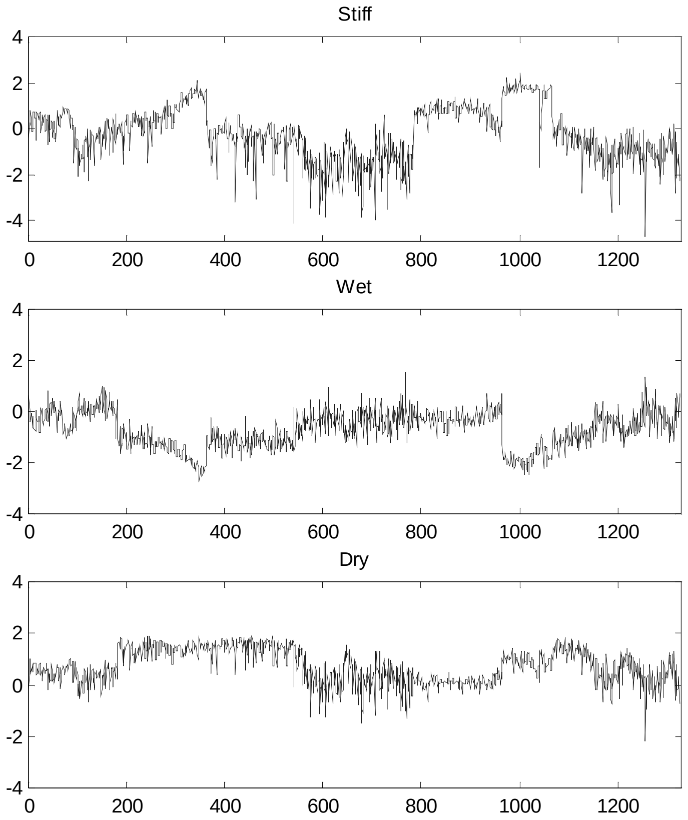

Fig 6. “Hand-made” features

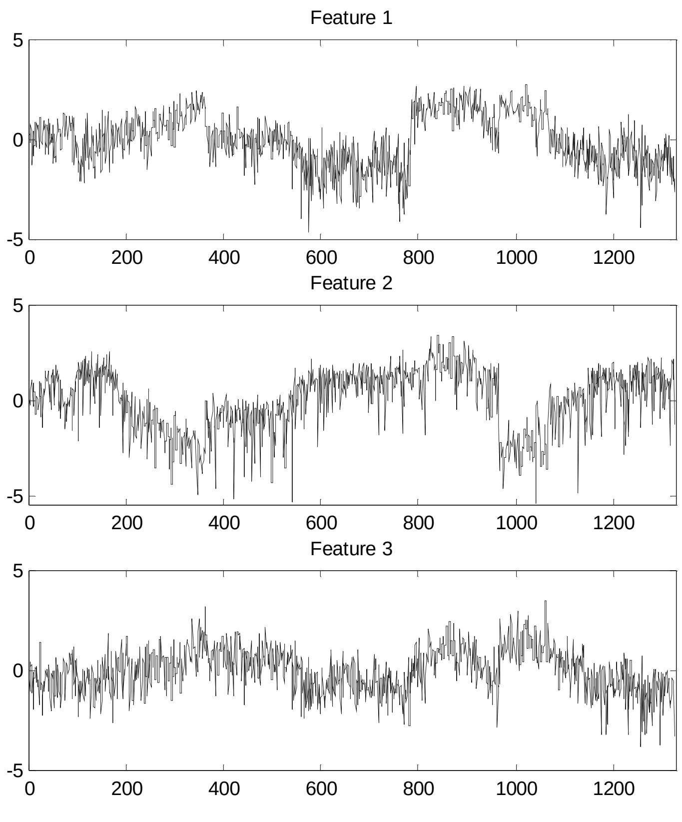

Fig 7. GGHA features

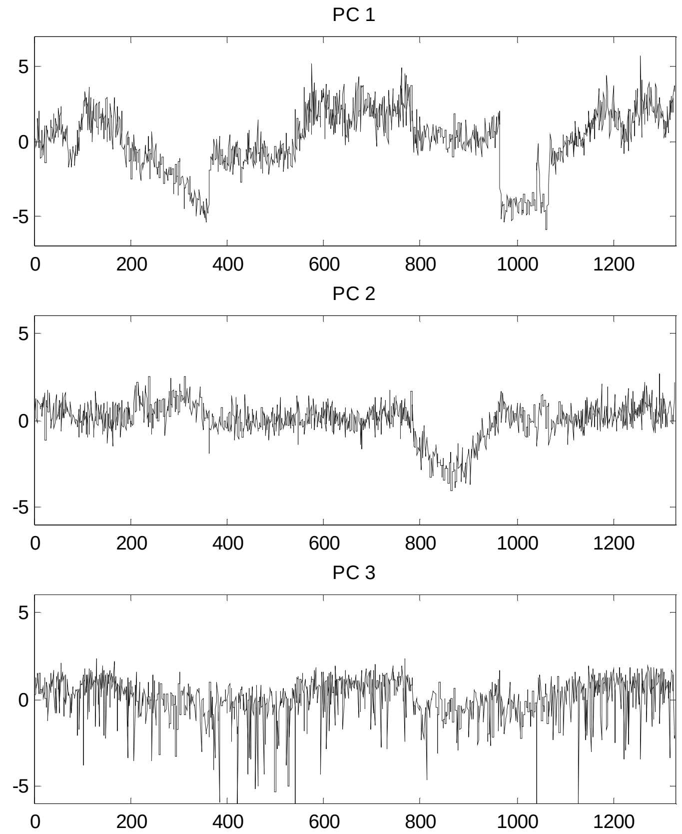

Fig 8. PCA "features"

	PC 1	PC 2	PC 3	GGHA 1	GGHA 2	GGHA 3	"Stiff"	"Wet"	"Dry"
PC 1	1	-0.0399	0.2887	-0.7749	0.6964	-0.7164	-0.7803	0.8185	-0.6336
PC 2	-0.399	1	0.0641	-0.3226	-0.4511	-0.5353	-0.1329	-0.1594	0.1922
PC 3	0.2887	0.0641	1	-0.2242	0.6707	0.0304	-0.0266	0.1686	0.0656
GGHA 1	-0.7749	-0.3226	-0.2242	1	-0.2488	0.7038	**0.8522**	-0.4611	0.3367
GGHA 2	0.6964	-0.4511	0.6707	-0.2488	1	-0.1312	-0.2735	**0.6446**	-0.4213
GGHA 3	-0.7164	-0.5353	0.0304	0.7038	-0.1312	1	0.578	-0.5616	**0.504**
"Stiff"	-0.7803	-0.1329	-0.0266	0.8522	-0.2735	0.578	1	-0.5305	0.4122
"Wet"	0.8185	-0.1594	0.1686	-0.4611	0.6446	-0.5616	-0.5305	1	-0.7618
"Dry"	-0.6336	0.1922	0.0656	0.3367	-0.4213	0.504	0.4122	-0.7618	1

Table 1. Correlations between the "classifications" of the test data. As is expected, "wet" and "dry" froth have high negative correlation; additionally, "stiff" and "dry" are positively correlated (lower right corner). It turns out that stiffness has strong positive correlation with GGHA feature #1, wetness correlates with feature #2, and dryness with feature #3 – these facts can be seen also in Figs. 1 and 2. Stiffness and wetness can be also categorized rather reliably by using the automatically extracted GGHA features, dryness detection using GGHA features is less trustworthy (looking at the numerical values, it seems that the froth was never specially dry). Note that principal components give no hints about the "correct" classification whatsoever. It seems that PC #1 has high correlation with all of the "mental features" simultaneously, giving no clue of how to distinguish between them!

4. CONCLUSIONS

There are many pragmatic problems when working with human experts. The experts cannot explain their judgements, but their pattern matching capability is immense … for example, it is typical that when some graph is shown to a domain-area expert, he (she) often can explain the curves in terms of underlying process phenomena – if the expert is co-operative, the results can be interpreted in a "positive" way even if there would not exist anything to be seen in reality. In this study this dilemma is avoided by consulting the experts before any experiments; the results are thereafter compared mathematically, in a "neutral" fashion.

The results here should still be regarded as preliminary. However, in [8], the model construction was carried out for different sets of input variables. Also in that case, similar results were obtained: there seems to exist at least some level of robustness and repeatability behind the hypotheses.

It needs to be noted that the categorization result is by no means unique: it reflects the relevant phenomena observed. The data set that was used for model construction in this case was collected during pre-planned experiments in the process: the inputs were varied in a wide range, so that the froth layer also reflected different kinds of operating conditions. Because of the special test circumstances, the data seemed to be rather free of "outliers", too. When the GGHA models were later constructed using some lower quality data, not all of the three nominal categories were found; it seems that "stiffness" feature is almost always detected, but sometimes the other features are needed to explain wide, spurious variations in the data perhaps due to measurement errors. Because these outliers are typically random, there is no correlation between successive samples in this sense; this random nature of outlier feature loading makes these uninteresting features rather easy to spot.

It turns out that the created feature model does not offer specially good latent variables for predicting the mineral concentrations in the froth. Now the only aim was to simulate the human perception – and humans are typically outperformed by statistical methods in this kind of numerical tasks.

Even though the results cannot be directly generalized, it seems that the school of expertise research could benefit from the new ideas. If the automatically extracted domain-oriented features may be easily understood by the domain-area experts, if they match their intuitive categories, they can hopefully be labelled with mnemonic names. A "common language" can perhaps be found, so that humans can communicate with each other using objective terminology, making the task of knowledge acquisition and maintenance easier.

ACKNOWLEDGEMENT

This research was carried out as a part of the EU funded project "Characterization of Flotation Froth Structure and Color by Machine Vision (ChaCo)".

REFERENCES

[1] Cipriano, A., Guarini, M., Vidal, R., Soto, A., Sepuldeva, C., Mery, D., and Briseno, H.: A Real Time Visual Sensor for Supervision of Flotation Cells. *Mineral Engineering*, **11** (6), 1998, pp. 489-499.

[2] Comon, P.: Independent Component Analysis – A New Concept? *Signal Processing*, **36**, 1994, pp. 287-314.

[3] Fisher, D.H., Pazzani, M.J., and Langley, P. (eds.): *Concept Formation: Knowledge and Experience in Unsupervised Learning.* Morgan Kaufmann Publishers, San Mateo, California, 1991.

[4] Hyötyniemi, H.: On Mental Images and "Computational Semantics". *Proceedings of the 8th Finnish Artificial Intelligence Conference STeP'98* (eds. Koikkalainen, P. and Puuronen, S.), Publications of the Finnish Artificial Intelligence Society (FAIS), Helsinki, 1998, pp. 199-208.

[5] Hyötyniemi, H. and Saariluoma, P.: Chess – Beyond the Rules. In *Games, Computers and People* (ed. Honkela, T.), Finnish Artificial Intelligence Society, 1999, pp. 100-112.

[6] Hyötyniemi, H.: From Intelligent Models to Smart Ones. *Proceedings of the Second International Conference on Intelligent Processing and Manufacturing of Materials,* Honolulu, Hawaii, July 10-15, 1999.

[7] Hyötyniemi, H.: *GGHA Toolbox for Matlab*. Helsinki University of Technology, Control Engineering Laboratory, Report 117, 1999.

[8] Hyötyniemi, H., Hasu, V., Hätönen, J., and Ylinen, R.: "Data Mining" for Mining Data. *Submitted to XXI International Mineral Processing Congress*, Rome, Italy, July 23-27, 2000.

[9] Hätönen, J.: *Image Analysis in Mineral Flotation*. M.Sc. Thesis, Helsinki University of Technology, Control Engineering Laboratory, Report 116, 1999.

[10] Kellogg, R.T.: *Cognitive Psychology*. Sage Publications, London, 1995.

[11] Lamberts, K. and Shanks, D. (eds.): *Knowledge, Concepts, and Categories.* MIT Press, Cambridge, Massachusetts, 1997.

[12] Laskowski, J.S. and Woodburn, E.T. (eds.): *Frothing in Flotation II*. Gordon and Breach Science Publishers, 1998.

[13] Moolman, D.W., Aldrich, C., van Deventer, J.S.J., and Stange, W.W.: The Classification of Froth Structures in Copper Flotation by Means of a Neural Net. *International Journal of Mineral Processing*, **43** (3-4), 1995, pp. 193-208.

[14] Olshausen, B.A. and Field, D.J.: Natural Image Statistics and Efficient Coding. *Proceedings of the Workshop on Information Theory and the Brain*, Stirling, Scotland September 4-5, 1995.

[15] Olshausen, B.A. and Field, D.J.: Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1? *Vision Research*, **37**, 1997, pp. 3311-3325.