Under the microscope

Robin Flowerdew and Barry Leventhal

Based on an article published in New Perspectives Magazine, September 1998

Despite their obvious success, no hard theory underlies today's geodemographic classifications. Robin Flowerdew and Barry Leventhal put one of today's products to the test and offer fuzzy geodemographics as a possible alternative.

Geodemographic classifications have been used in the UK for the last twenty years and have contributed to countless targeting projects. They have helped in planning site locations for numerous retailers – the total value of these investment decisions alone must surely run into billions of pounds. Yet where is the proof that geodemographics can actually predict how markets are geographically distributed, and which classification approach works best?

Rule of thumb

There is no formal proof and no "theory of geodemographics" either, only the concept that "birds of a feather flock together". All the evidence is empirical and the practical results tend to stay within the companies who have tried the technology. The systems are used simply because they do work and have become established within the companies. Doubtless there are instances where they fail and other methods perform better, but these cases are never publicised – after all, why pass this knowledge on to your competitors?

Both the Market Research Society's Census Interest Group (CIG) and the Group Market Research department at Whitbread plc felt these were important issues. Whitbread is one of the UK's major retailers, running a collection of retail chains in the pub and restaurant sectors, and Group Market Research has been a seriously heavy geodemographics user since the birth of the industry.

A working party was formed in 1995 to examine the questions and try to find some answers.

The party soon grew to include representatives from BMRB International, who operate the Target Group Index – the TGI is arguably the most widely used source of "geodem cross tabs" – and also the two leading census agencies, CACI and Experian.

The working party started by looking at geodemographic profiles of product usage on the TGI and saw that the discrimination was impressive, with some products, such as the classic dishwasher ownership example, showing up-market bias and others, such as cigarette smoking, showing down-market bias. However, we wanted to test whether differences in area profiles really predict differences in consumption at a neighbourhood level and the only way to do this was to measure actual consumption rates for neighbourhoods within a town, and compare the observed behaviour with the geodemographic predictions.

The Luton survey

The town selected was Luton, an area well known to the Whitbread representatives; this local knowledge was useful for sense-checking the results. The Luton local authority district contains 340 neighbourhoods in the form of census enumeration districts (EDs), the units of geography for the 1991 Census.

The test was structured using ACORN, which classifies EDs into 57 different types on the basis of 1991 Census data. Similar results could doubtless have been produced using other neighbourhood classifications and the results were post-analysed using MOSAIC, a postcode-based discriminator, with corresponding findings.

Three ACORN types were chosen for the test corresponding to high, medium and low involvement in a wide range of consumer markets. A sample of nine wards was selected containing the target types and, within these, two EDs were chosen per ward. We thus ended up with a sample of six EDs for each of the three target ACORN types – eighteen in all.

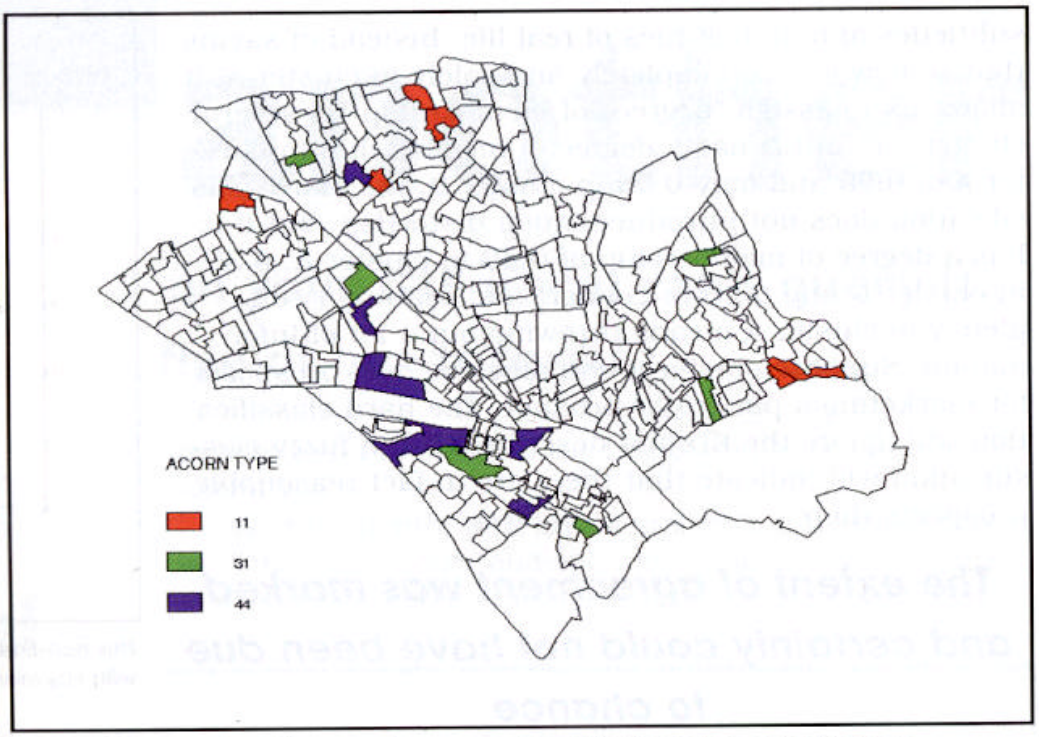

This map shows where the sample enumeration districts were located in Luton.

A sample of addresses was drawn for each ED, using the Ordnance Survey's AddressPoint product to ensure that all of the addresses were actually located within the designated neighbourhood.

Next came the fieldwork. BMRB interviewers conducted the survey during September 1996, interviewing one adult per household from the list of addresses. The questionnaire covered demographics, product consumption and media usage, and in some ways was equivalent to a mini-TGI survey. The field workers repeatedly revisited the selected addresses to make sure we got the target response, and interviews were achieved with 86 per cent of the eligible homes, 867 in all. With such a high response rate, we can be confident that the results are representative of the actual behaviour of the population within each neighbourhood.

The initial analysis looked at the pattern of consumption rates by neighbourhood types, comparing the survey results with predictions from the TGI. The extent of agreement was marked and certainly could not have been due to chance.

With a collection of measurements on each respondent, we really needed an approach that accounted for the relationships between the consumption patterns, so multivariate analysis seemed highly appropriate.

Initial clustering was conducted using the product consumption profiles to form a classification of the eighteen EDs. If the ACORN system is performing well, the EDs' consumption patterns should form clusters with the same composition as the ACORN types. When compared with their ACORN types, there was agreement in fifteen cases and disagreement in three: a good but not perfect performance. At this stage, we wondered how we could explain the behaviour of the three that didn't agree (EDs 5, 7 and 16 highlighted below) — was it due to the classification approach, or change since the 1991 census, or some other factor?

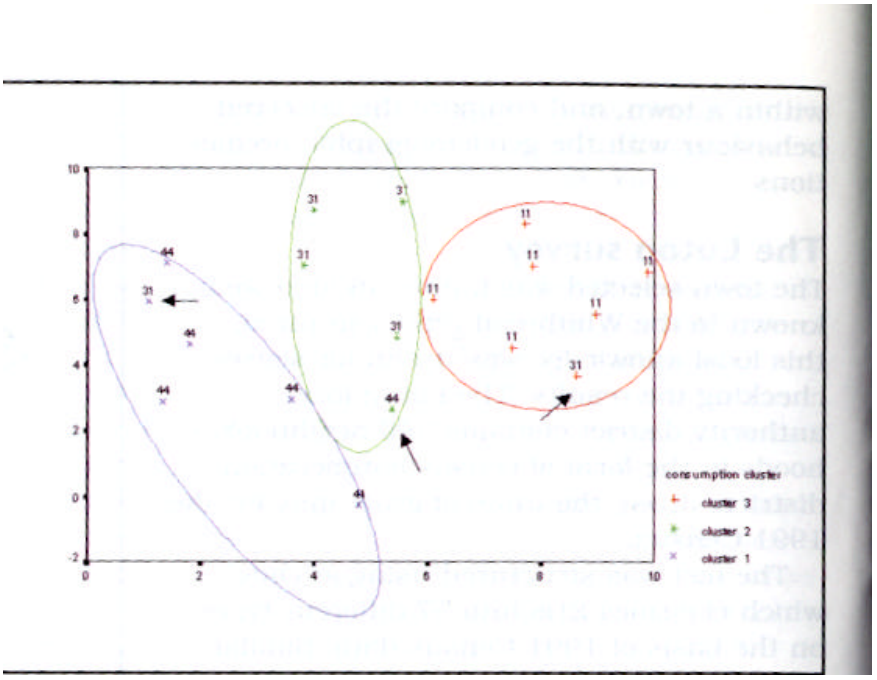

This non-linear map shows the results of clustering the consumption data, with EDs identified by their ACORN class.

Having applied the conventional "hard" cluster analysis technique, which assigns each neighbourhood absolutely to the closest available type, we decided (with the help of Zhiqiang Feng of Lancaster University Geography Department) to test an alternative "fuzzy" approach.

Fuzzy classifications

We normally think of a classification as an "either/or" system. If something is in cluster A, it cannot be in cluster B or any other cluster. Sometimes, this approach may be a bit too clear-cut; a place may be quite like the places in cluster A, but also be a bit like the ones in cluster B, and perhaps have similarities to places in clusters C and D as well. To put it clearly and unambiguously into cluster A may tell only part of the story.

A fuzzy classification explicitly allows us to tell the whole story without discarding the subtleties and ambiguities of real life. Instead of saying that somewhere is completely and solely in cluster A, it allows us to assign "degrees of membership" in several clusters. If an ED has a degree of membership in cluster A of 0.95 and only 0.05 in cluster B, the 'hard' classification does not introduce much distortion, but if it has a degree of membership of 0.40 in cluster A, 0.35 in cluster B and 0.25 in cluster C, allocating it completely to cluster A means throwing out a lot of information. Suppose cluster B contains the best prospects for marketing a particular product. The hard classification will ignore the ED just described, but a fuzzy classification will indicate that there are in fact reasonable prospects there.

One reason for the mismatches observed in the Luton survey between the ACORN system and the consumption clusters may be that the use of a hard classification of the consumption data forces each ED into a rigid mould, assigning it wholly to one ACORN type when it may actually have similarities to others. A fuzzy classification can identify if this is the case.

The results from the fuzzy classification of the consumer data show whether the EDs in the Luton survey fall clearly into a particular cluster or are intermediate between two or more. Some EDs (6 and 11 for example) have degrees of membership in their cluster of over 0.9, which is pretty

clear. Others, like ED 13, have medium degrees of membership in two clusters - it is most similar to ACORN type 31 (0.485) but is also like ACORN type 11 (0.358). However the three outlier EDs identified in the hard classification all have fairly strong links to a single cluster - unfortunately the wrong one! So fuzzy classification has produced some worthwhile additional information but has not sorted out our three problem EDs.

Fuzzy ACORN?

The next step in our attempt to find out why these three EDs don't fit our model is to try the fuzzy classification method out on devising the geodemographic categories themselves. Using census data for the 18 Luton EDs, therefore, we tried to produce a "fuzzy ACORN". We didn't have the precise list of variables used in ACORN but we did our best to replicate it as well as we could. Our fuzzy classification was not very different from ACORN, the main difference being that ED 5 (one of the three misfits) turns out to have a stronger degree of membership in the equivalent of ACORN type 11 (0.693), matching its product consumption patterns, instead of type 31 to which ACORN assigns it.

This still leaves two outlier EDs unaccounted for. Subsequent research showed that ED 16 had undergone substantial change in its character between the 1991 Census and the 1996 survey, accounting for its wrong assignment in both the fuzzy and hard classifications. This was established through comparing those survey respondents who had moved in since 1991 and those who had been there for more than five years. The new residents were significantly different in terms of age, education level and employment, though similar factors could not account for the final rogue ED (7).

Our final idea was that the general-purpose ACORN classification, and our fuzzy version of it, might not be targeted specifically enough for the consumption variables we were interested in. So we created a special-purpose fuzzy geodemographic classification, giving higher weights to those census variables which were most closely related to the specific consumption variables used. This time ED 7 came out with high degrees of membership in two clusters, and therefore an intermediate status between two clusters, one equivalent to its ACORN type and one matching its

consumption patterns. All three anomalies are now accounted for.

Evaluating the results

The research gives us a proper basis on which to evaluate the success of geodemographic targeting and to compare different methods. ACORN passed the test well but the wrong assignment of three EDs means that it isn't perfect. Fuzzy classifications do rather better. To explain our survey data completely, we need to consider how places have changed since the census, and to construct a special-purpose classification, emphasising those census variables known to be good for predicting consumption of the specific goods in our survey.
So does fuzzy work best? We think it probably does, not just because our fuzzy classification fitted the Luton test data slightly better than ACORN, but because the concept of degrees of membership of a cluster offers a more flexible and sensitive approach to targeting.

Robin Flowerdew is now with the School of Geography and Geosciences at the University of St Andrews. Barry Leventhal is now with Teradata, a division of NCR, and is chair of the MRS Census & Geodemographics Group.

Our thanks to James Lawson, formerly editor of New Perspectives and currently editor of Database Marketing (www.dmarket.co.uk) for his help in sourcing this article.