

GDI IMPULS

Wissensmagazin für Wirtschaft, Gesellschaft, Handel
Herbst 2005

Andreas Weigend

«We don't know, but
we can measure it.»»

Reprint from GDI Impuls, Autumn 2005: «Wege aus dem Billigsumpf»



ANDREAS WEIGEND is an experimental physicist, an expert in data mining, Amazon.com's former chief scientist and a consultant to international companies in the USA, Europe and Asia on strategic and analytical projects. He has published more than one hundred scientific papers and co-authored six books. Until 1999, he was full-time faculty at New York University's Stern School of Business. Previously, he taught as assistant professor at the University of Colorado (Boulder). He currently teaches at Stanford University, at the University of Washington and at two of the most renowned Universities of China: the Shanghai Jiao Tong University and the China Europe International Business School. 1999 Weigend co-founded MoodLogic in Silicon Valley. Their main product was awarded best software for the organisation of music. Weigend lives in San Francisco, Shanghai, and on weigend.com.

«We don't know, but we can measure it.»

As Amazon.com's chief scientist, Dr. Andreas Weigend developed a huge database of customer data, that brilliantly blended product strategy, usability, market research, and marketing. This interview reveals this thought leader's vision for the next level of customer-centric companies.

This interview appeared originally in the Autumn 2005 issue of IMPULS in German. The interviewer was Stefan Kaiser of GDI.

«Where should we have the shopping cart – left or right?» Whoever talks to Andreas Weigend, the former chief scientist at Amazon.com who now works as a consultant, must be ready to question every detail. Weigend's matter-of-fact answer to this question is: «We don't know.» As a leading data mining specialist, however, he has the answer: «But we can measure it.»

And that is exactly what Amazon.com does. All the time. In the above example, Amazon.com presented half of its customers with a website that differed in only one detail – the side on which the cart was located – and then compared the purchasing behaviour of both groups of customers.

A POWERFUL METHODOLOGY

This way of designing experiments and rigorously validating the results has enabled internet stars like Amazon.com, Google and Yahoo to maximise their profits and provide a superb customer experience. A quick, inexpensive experiment will often tell you a lot more than focus groups or elaborate academic studies. And with more telling results: preferences revealed in the real environment often differ from stated preferences

in an abstract setting. And the results are a lot faster in coming: rather than months or weeks, the internet answers the question posed in the experiment in days or hours or sometimes even minutes, thus accelerating evolution by orders of magnitude. And since the person carrying out the experiment needs to provide the group with differing situations (such as left and right), actionability is intrinsic rather than guesswork based on insights.

Fast feedback from experiments is one of the advantages of the online world. Hundreds of experiments can be carried out on a website at the same time. And if the groups are run in parallel, many variables that are hard to control or observe, such as the weather, average out since customers are assigned at random to one of the groups. A complete record of the customer's actions is another advantage when building predictive models of customer behaviour.

One of the more surprising insights was just how predictable people are. Decisions a customer believes were made spontaneously or after much deliberation are often more predictable than the customer would have ever guessed. The

clicks that millions of customers generate day after day provide a rich basis from which to establish behavioural patterns: «We often anticipate what people need before they realize it themselves.» Weigend comments.

THE TWO PARADIGM SHIFTS IN MARKETING

Marketing is currently experiencing a paradigm shift. Approaches such as traditional segmentation, merchandising and branding are giving way to rigorous measurement, experimentation, and data mining. And the effect on the bottom line? Amazon.com's recommendations account for 10 percent to 20 percent of its sales. And Toys 'R' Us reports that computer-generated recommendations have doubled the number of sales per click on their website.

But this is only the first of the two paradigm shifts. The potentially even more important second shift focuses on providing what are known as architectures of participation that allow users to contribute their knowledge. An early example here are the millions of book and product reviews at Amazon.com. More recent examples include (i) Wikipedia (www.wikipedia.org) – a collection of



Live market research

Live experiments give rapid answers in market research. Amazon.com presents its customers with two versions of its website, differing from each other only in one navigational aspect, the position of the check-out box. The version on the right generates an increase in sales of between one and two percent.

articles on essentially any topic, written collaboratively by millions of contributors, (ii) blogs – diary-like weblogs currently being written by between 10 and 20 million individuals and read regularly by about one tenth of internet users, and (iii) tags – words surfers use to label a website they want to remember and are willing to share with the rest of the world. Sometimes called social software, these are examples of platforms that allow users to contribute to the organisation of information. They are active creators of meaning, not slaves to the links provided by the pages.

In the new world of marketing, customers are no longer targets for ‘push’ advertising messages – even if they are selectively tailored, but are decision makers who ‘pull’, who express themselves as individuals, who enjoy the discovery process, and are interested in the behaviour of other people. Thought leaders such as Andreas Weigend see in this dramatic change of perspective the central objective for the further development of marketing. He is working towards ensuring that data mining has an important role in this new environment.

Dr. Weigend, how does a PhD physicist from Stanford end up observing people shopping?

Andreas Weigend: Well, once a physicist, always a physicist! In the mid-80s I worked at CERN in Geneva, then at Stanford’s Linear Accelerator Center SLAC, where I worked on identifying patterns within data. After graduating, many of my fellow physics graduate students ended up on Wall Street, and I, too, created trading models for Goldman Sachs, Morgan Stanley, J.P. Mor-

gan and others while I was a professor at New York University. As in physics, we built empirical models based on data and identified patterns in these data - of course, not in the information left by elementary particles, but those left by the markets. And in that regard, my

«It all began with the philosophy of Amazon founder Jeff Bezos, that we never throw any data away.»

move over to e-business meant study of just one more data source, the information left by individuals. I was identifying patterns and building predictive models for human behaviour.

The methods interested you more than the customers themselves?

Good research requires three elements: first, the data; second, the appropriate methodology, and third, domain expertise - an understanding of other disciplines such as behavioural economics. If a person is just a statistician, and only interested in methods, but does not know where to look, or has no data, it will be reflected in the results.

At Amazon.com you developed pattern recognition to perfection. Amazon is considered to be the benchmark for individual customer contact and service. How did this ‘customer-centric approach’ develop?

It all began with the philosophy of Amazon founder Jeff Bezos, that we never throw any data away - everything that could be gathered, would be gathered.

On Christmas Eve that information can amount to a terabyte of raw data. It is also equally important to hire people who have IT capabilities and are interested in making something out of this data and putting something together for the customer. These two ingredients are

what led to the hundreds of features that you see on the Amazon.com website.

One highlight is the individual recommendation system. The automatically generated tips are usually astonishingly accurate.

Some recommendation systems are based on the click or buy history of customers, and some attempt to assess the current situation. Amazon.com is known by the public for the former, although the data for most customers is not extensive at all. In terms of cross-selling, in particular, it is very difficult to recommend a frying pan on the basis of a novel someone has purchased. This kind of history is inadequate for advertising purposes and other hoped-for applications.

‘Item-by-item’ recommendations are much more important. The idea is simple: we set up a very large matrix in which every single product appears in relation to every other product. This matrix gives information on the purchasing behaviour of the customer in regard to multiple purchases - for ex-

ample, when shirts and ties are bought at the same time. The coincidences are averaged out statistically, and the signal becomes apparent. This and other matrices - for example, on click behaviour - capture the current situation of the customer. If we now add to this the behaviour of every customer in regard to a particular product that someone has just placed in his or her shopping basket, the results are frequently astonishing.

It is easier to predict what else a customer might be interested in if he has placed a certain item in his basket - using data on what other customers in the same situation have selected - than to predict what he might buy next year. The few search terms and clicks used on any given day usually say much more about the intentions of a customer than any extensive customer history.

You really need only a few clicks to determine what a customer in a particular situation will do next?

If we have all the past data on every customer: yes. An extreme example of this are Yahoo and Google ads, where on the basis of one single search term, the system will show ads whose relevance is frequently very accurate.

So the decisive change in online marketing was the move from past-behaviour analysis to situation analysis?

That's right. But it is not really anything new: when I want to buy myself a suit, a good salesperson assesses the situation he sees in front of him and does not ask me what kind of suit I was given for my first communion.

You say marketing people should not settle for any old study, but should concentrate on possible actions early on.

Data mining is often unsuccessful because people do not know what to do with the insights gained. In contrast, if I know what types of action I can control - the products I can offer together,

period of time - you know that everyone sees a different Amazon site from his or her neighbour. This consistency over time is very important if we are trying to measure long-term effects like customer satisfaction. Second, customers leave tracks in the internet: their click behaviour reflects their decision-making

«Marketing specialists pose questions and let the customer answer them in the same way that physicists pose questions that nature answers.»

the discounts I can give, the e-mails I can send out, or the special offers I can push, that is when data mining becomes unbelievably strong. If you just look at the data piecemeal, you will find patterns, but most of the time the reaction is: «We knew that» or «What can we do with the information?» In marketing today it is very important to tackle the problem from both sides concurrently: analysis and actions.

You also do experiments. What advantages does the internet world offer you?

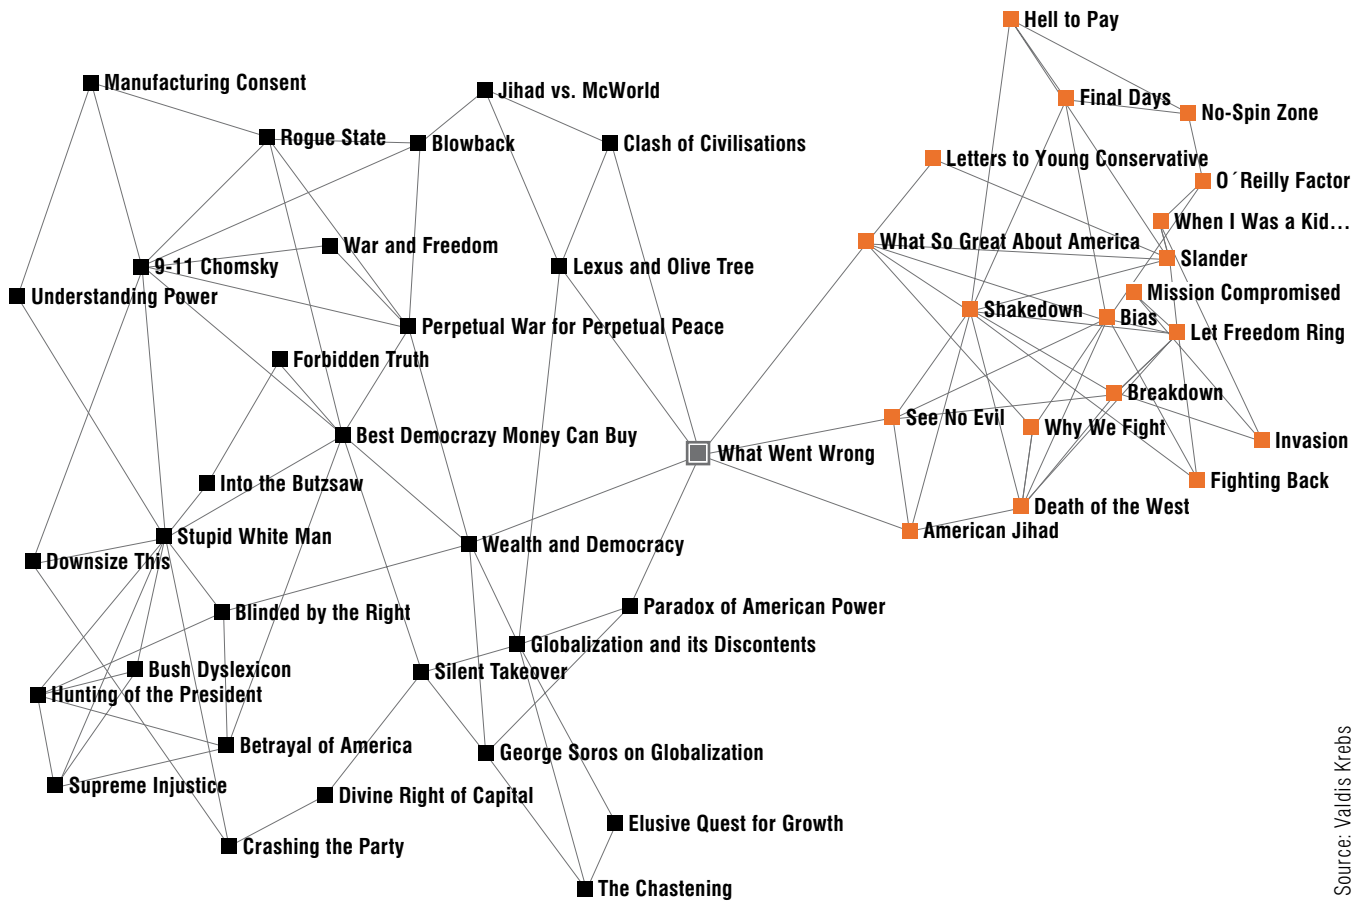
Experiments make us consider the conditions that we can change. We measure reactions to two alternatives, both of which are workable. So marketing specialists pose questions and let the customer answer them in the same way that physicists pose questions that nature answers. The online world makes this very simple: First, every user can be shown a particular page layout over a

process and gives very good indications of what they find important. A third important characteristic of the internet is rapid feedback for the person carrying out the experiment: we can get answers to our questions in the shortest space of time. Amazon.com sends out a thousand pages per second; when we want to test a new idea on an experimental basis, we will get back statistically significant information on how it affected customer behaviour on the same day.

You throw out bait to find out how you can reel in customers better ...

... Experiments are more than bait. Experiments are a philosophy and, in science, they represent the only way to find out about the world. They are also essential drivers of innovation. Because there are many people at Amazon.com who are in a position to test their ideas through experimentation, it is not just a few market researchers, but a lot of IT people who also take part and

BOOKS SALES AND POLITICAL OPINION



This data from Amazon.com perfectly mirrors the political situation in the USA. ■ Black Boxes reflect books more critical of George W. Bush. ■ Red Boxes reflect books from the pro-Bush camp. The lines show common buyers. Both sides are linked only by one single title: «What Went Wrong?».

try out their ideas, both large and small.

This is how innovation is built into the system. We empower intelligent people, not only - as in many other companies - to carry out tasks given to them by marketing, but also to develop and apply their own creativity.

Can you cite an example?

A friend of mine at Amazon.com ran an experiment to get as many customers as

possible to apply for an Amazon.com credit card. The difference between: «Get thirty dollars off your next purchase» and «Pay only ten dollars today» if you have forty dollars or more in your cart is astounding. He tried out the two versions, and the result was that four times as many people applied for a card in response to the second approach.

These 'framing' effects in decision making discovered by Nobel Prize Win-

ners Daniel Kahneman and Amos Twersky ...

...- two leading representatives of behavioural economics -

... are also observed in the free for all of commerce in the form of multi-million dollar sales revenues.

You coined the wonderful phrase «quero ergo sum» - I search, therefore I am. If man is the sum of his search queries, or if these queries represent

the essence of the person, then it should certainly be possible to assess customers on the basis of their search behaviour.

The most telling concentration of our lives is represented by our search queries in the web, because I only search for things that I do not yet know, but want to know. The fact that Yahoo or Google make it possible for us to find useful information in spite of the enormous mountains of misinformation is an enormous achievement. Our search expresses what interests us right now. And, in essence, that is what is responsible for our reaction to directly related advertising, and – in Google, for example – to sponsored links.

A new form of search is based on what are known as ‘tags’ – labels that visitors give to websites they have visited. What do you think of this phenomenon?

Tags are fascinating. www.del.icio.us allows us to attach words to websites that remind the user of the website. But the real value comes when I follow not only my own labels, but those of millions of other users. That means that I can follow users from a website and see what else they found interesting. I can also follow a tag and see what else on the web comes under the same category. The fast-feedback loop is important here, too. I think of a descriptive word, enter it, and am immediately shown websites that have been tagged by others in the same way. The helps me find out if the tag is the right one, or if a different tag would be better.

One additional point can be illustrated by the four phases of the internet. In its first phase, information was simply

uploaded onto the net, but users had to know where it was. In its second phase, authors showed the way to information using hyperlinks. In the third phase, readers themselves pointed the way. Hyperlinks are relationships created and positioned by authors while tags express relationships created by readers.

Google made a significant contribution to the second stage of the internet in that it created a link structure based on relevance – without which we would drown in a sea of results. Now readers are responsible for producing an even greater flood of information. In the fourth phase of the internet, it is time to apply Google’s success to web diaries; to blogs, phlogs, vlogs, and podcasts. A new blog is created every second – how can we find out what is of interest?

The algorithms being developed by firms like Rojo.com use person-to-person links, as well as hyperlinks, that users record in social networks (LinkedIn, openBC). They also have ‘attention data’ that describes the click and read behaviour of users. This fourth phase will allow us to discover things that we never knew we could look for. I believe this will play a central role in marketing.

That gives us a new way of creating recommendations ...

Yes...

...If we apply these thoughts to marketing this could represent another ‘Copernican Turn’. Instead of bombing the consumer with dull mass-marketing messages, customers themselves will take an active role.

The Copernican Turn is the act of tur-

ning away from a frequently undesirable push – like when someone sends me an ad via a text message for which I actually have to pay a small charge – to pull behaviour. I want to know what interests those other people who created the tags. The process of discovery promoted by these ‘social bookmarks’ is what is important here. The Turn is marked by a fascination for following others, getting to know them, and sometimes even developing an emotional bond to people about whom you know nothing other than you like the way they find information, and you also find the websites that they have tagged interesting.

Does that not also signify a change for data mining? Previously you analysed people based on their search and click behaviour, now you almost have to start up a dialogue with them.

Technology like tags makes communication between customers possible, instead of just having a one-way conversation between the provider and customer.

Thanks to communication and to fast-feedback loops, we can now focus and amplify the collective intelligence of people. I get feedback on what my tag means and I adapt it accordingly. This puts us in a world where we can follow up on interests and discover them for ourselves, rather than one in which we are simply told to follow the superhighway.

The area where data mining or machine learning plays an increasingly important role is in sorting and ordering search results and recommendations. The skill is in showing those pages, products and people first that have the highest probability of being relevant to a

user in a given situation. This question of relevance is now also coming up for tags – which can call up thousands of pages – for products that appear in reviews with thousands of conversations as well as for social networking and on-line dating sites.

What elements are shaping this new world?

For tagging and partnership sites, where the people themselves are the ‘product’ that must be sorted according to rele-

from those who typically come to this site. There is often no offline alternative available. The phenomenon of discoveries made by individuals who divulge only a little information about themselves – that is, however, very detailed – is specific to the online world.

But still the net has created structures that facilitate social connections.

Today many young people assure themselves that they are in contact with their kind of people by texting or instant mes-

that he spends on himself and his family with a particular company. For some people the customer network value is much higher: they influence others to spend overall higher amounts with a company than they ever could on their own. Such people enjoy a much higher level of trust than do marketing people or advertising displays. It is worth a lot to companies to find these people and look after them.

The concept of opinion leadership is well known in classical marketing. What is the difference here?

People try to find typical opinion leaders in the offline world – say a golfer who shows off the watch he wears. In the online world we break that down to each individual, so that the algorithms assign a customer network value to each of the 50 million customers of a large company. This is completely new. The salesperson in a traditional store does not know what my customer network value is. He does not know to how many other people I will recommend his shop. On the other hand, with an online retailer, this information will be part of my personal details in the individual reputation management system.

...which should also be worth money to the customer.

Absolutely.

Development here is strongly driven by the mobile telecommunications industry. Where do you see the next areas of interest for data mining?

Developing technology and algorithms that address the actual situation in which people find themselves. That

«In the times before Gutenberg, women working on market stalls already had a feel for the relevant.»

vance, we see three important elements: interaction, status modification while absent, and random rewards. This mix easily leads to a dangerous dependence, such as that associated with online stock trading in the late 90s, or from online gaming nowadays. In all of these cases an nebulously defined world beckons, offering positive rewards that are allocated on a random basis – and every now and again you win the jackpot. This is a point where society will change radically, and marketing has to decide how it will respond to the change.

Is there any online-offline transfer?

When the airline ‘Swiss’ runs an advertising spot that refers to their website just before the ‘10 to 10’ news, we can measure how many additional people go to the site. We can also observe how the behaviour of these people differs

saging – much like apes who show their affection by calling out and stroking each other. This does not usually transmit any important information other than: ‘What’s up?’ ‘Not much.’ Given that communication costs all over the world have reduced so dramatically, many more multilayered networks have come into being than was previously the case in the era of the tiny mountain village.

Marketing also tries to use these structures. Research is currently ongoing into customer network value that addresses how many purchases I can generate through my direct influence on others. What is that about?

Individual customer networks are a very promising area of research. The important factor here is that the lifetime value of a customer is not just that amount

means ‘situation-aware communications’ – and marketing is, after all, a form of communication – to predict my reaction in a given situation based on past behaviour. Wouldn’t it be great if your mobile phone could learn when you want to be disturbed and when not? The question is not ‘Should we send a text message or not?’ but rather ‘When is it the right thing to do?’ Everyone would have peace and quiet when in an important meeting, but receive important information when waiting at the airport. There is still quite a lot of work to be done here.

This premise could easily be extended across all communication channels. There should be less focus on the situation and more on its respective relevance to customers.

Before the Gutenberg process, women working on market stalls already had a feel for what was necessary. They sensed whether a customer was happy to chat or whether he or she simply wanted to buy a head of lettuce. Then there was a phase at the turn of the millennium when the interactive aspect was ignored by one-way communication media. Only slowly did feedback channels teach us how rich and personal things had been in the pre-Gutenberg era. The difference is that we no longer communicate only with the 50 people in our village, but with the 50 million people on the web. Along with customer network value and situation awareness, this two-way communication represents an important area of research.

Is the question of relevance still topical? It has taken on increasing importance

over the last few years. Ten years ago Colin Harrison, of the IBM Research Laboratory at Rueschlikon, spoke in a brilliant lecture about ‘The Next Big Thing,’ – people would develop the need to access information that is of personal interest to them. Since then our knowledge on the relevance of data has raised our expectations exponentially. Previously, if a trip to the university library resulted in a book that touched on the area of interest in question, it was considered to be successful. Today, given that on Amazon.com text searches can be made that even include the content of a book, much more precise search results are expected. Our expectations in regard to personal data 10 years ago now represent the benchmark for expectations in regard to the relevance of texts and content. On dating sites – a market that is, by the way, worth billions – people expect that having entered just a few criteria, they will be matched with people who are available and interested and even live in the same neighbourhood.

Do you believe that we will get there? The importance lies in getting ‘real’ hits. Fake relevance would turn off consumers very quickly.

Who measures ‘relevance’? Relevance is always subjective. One can ask people if they got what they had expected. But people who are looking for something do not know necessarily know what it is. The problem first emerged when search queries resulted in large numbers of products, texts, books and people. Previously, relevance had not posed a problem. But when I think of the many millions of blogs that exist, I suggest

that we do not call it a problem of relevance, but more of a problem of discovery. <

Gottlieb Duttweiler Institute

GDI Impuls

Langhaldenstr. 21

CH – 8803 Rueschlikon/Zurich

Switzerland

www.gdi-impuls.ch