

Intelligent profiles and segments equals pure power for business: Combining profiles, segment and predictive analytics

Received (in revised form): 4th November, 2015



Andrea Ahlemeyer-Stubbe

is Director of Strategic Analytics at HackerAgency München GmbH and the author of 'A Practical Guide to Data Mining for Business and Industry' (Wiley). Upon receiving her Master's degree in statistics from the University of Dortmund, Ms Ahlemeyer-Stubbe formed a consulting firm, offering customised professional services to her clients. She now leads HackerAgency Munich's analytics team, drawing on the wealth of experience gained from her 20 years in the industry, specifically in the areas of data mining, data warehousing, database marketing, CRM, big data and social CRM. Ms Ahlemeyer-Stubbe is a frequent lecturer at several universities, as well as a speaker at professional conferences. She was President of European Network Business and Industrial Statistics (ENBIS) from 2007 to 2009 and is currently Co-chair of the Industrial Conference on Data Mining (ICDM).

HackerAgency, Konrad-Zuse-Platz 11, 81829 Munich, Germany
E-mail: andrea.ahlemeyer-stubbe@HAL2L.com



Stephan Horvath

is Global CMO for HackerAgency, and as such is adept at combining customer insights with marketing technologies to bring tangible, measurable benefits to clients' business. He brings 20 years of experience in marketing and sales consulting for several industries, and holds a Master's degree in management science and computer science from the Technical University of Stuttgart. Mr Horvath leads teams to develop successful relationship marketing strategies in highly competitive market environments. His expertise spans the entire customer lifecycle: generating qualified leads, increasing the sales conversion rate of those leads and turning them into profitable, loyal customers. He drives lead and loyalty programmes for automotive, airline, finance and retail clients in Europe, North America, the Middle East and Asia. Mr Horvath has worked with leading global brands, including IBM, T-Mobile, Vodafone, Lufthansa, Coca-Cola, Hertz, Citi, santander and Volkswagen.

HackerAgency, Konrad-Zuse-Platz 11, 81829 Munich, Germany
Tel: +49 172 7111003; E-mail: stephan.horvath@hal2l.com

Abstract Business is excellent. Customers are more than loyal; they are in love with the brand. Marketing campaigns regularly bring good results and new customers. Innovative products precisely meet current needs. Who would not like to say such things about their company? The way to get to this nirvana is through thoughtful consolidation and analysis of all relevant information about customers and prospects in order to calculate profiles and segments — and to extrapolate them into the future. The general use of audience segments is not new, and marketing profiles have existed since before data grew big. But by collecting and using online and offline behavioural data, location data and touchpoint data, in combination with the power of predictive modelling, we can provide insight-driven, individualised communication and interactions with customers and prospects. This provides indispensable fuel for daily business and drives client success. This paper will outline the framework for these newly developed and successfully implemented methods and will describe some of the business opportunities they can empower.

KEYWORDS: intelligent profiles, intelligent segments, automated predictive models, predictive modelling, unsupervised learning, clustering, imputation

INTRODUCTION

There has never been as much data from prospects and customers, from so many possible sources, as there is today.

While data protection laws may be more or less restrictive in different countries, some industries function under tight regulations and some clients are very sensitive about privacy issues, it is still true that the world of accessible data has changed and grown into a universe. While still under discussion at the time of writing, the US–European Union Safe Harbor Framework will likely allow US companies that self-certify to transfer personal data outside the EU, consistent with the EU Data Protection Directive — adding more data to the mix.

In addition to the inventory data that every company stores in its customer database (name, gender, address, customer ID number, orders, purchases, requests, customer service enquiries, payment history, etc), we can, and do, amass much more information. For example, we collect data about the online activities of customers and prospects in social media, their registrations in online forms, their usage data from websites and apps, their click behaviour and interactions, their completed or aborted purchase transactions and their preferred devices and browsers.

Offline data also arise from visits to the retail location, from responses to offline campaigns, and from regional peculiarities and market environments, to name just a few.

Offline campaigns are marked with codes. Most point-of-sale systems measure activities there, while mobile devices provide other geolocation data.

Online and social media marketing activities get unique URLs, or cookies, to help track the origin of user actions. Information about consumers' preferred communication and information channels are obtained. It does make a difference whether visitors have found their way to

the company's website via the Facebook fan page, the company's profile on Xing or LinkedIn, a specific mobile ad, the company's e-newsletter, the corporate video on YouTube or the QR code in a magazine display ad.

Through monitoring tools and services, as well as application interfaces (API interfaces to web or social services, which can be used to read personal data), we can gain other valuable information about consumers' education, lifestyle, peer group information,¹ preferred brands, exposure to areas of interest and more.

There are two crucial points to keep in mind:

- A company has to respect national laws. In many countries an active declaration of agreement by consumers — both prospects and customers — is required for the storage and use of their data for marketing and targeting purposes. This can be a communication challenge, but it seems to be getting easier. People have become more willing to pay for information or entertainment with their data.
- The more elusive point is the issue of data matching and how to add data from different sources to a single person's record. As a rule of thumb, in the early stages of the consumer lifecycle, matching points are sparse and may not lead back to a single person; but they may lead back to a group of people who have something in common. As the relationship grows, more contact data and behavioural data become available for identification and matching. The goal is to try to collect these data as soon as possible.

But all of this information attached to a person is like crude oil — until it has been filtered, cleaned, analysed and aggregated, it cannot provide its full effect as fuel for business. Technologies and analyses are therefore not ends in themselves. They are

used to provide the best added value for consumers and businesses.

Consumers rightly expect businesses to communicate and interact with them as individually as possible. They expect relevant offers and quickly feel bothered by irrelevant advertising that offers no particular benefit. Only those organisations that know their customers and prospects well can build and maintain a mutually profitable relationship.

PROFILES

Using descriptive, exploratory analysis, unsupervised learning techniques and predictive modelling, we can create individual, precise and informative profiles from consumer data. These profiles consist of social, self-reported and observed behaviours, blended with classical CRM database information, which allow insights into how individuals act and connect in real life.

What a profile contains depends on the available data and specific objectives of each company. Apart from the usual CRM data, for example, a hotel chain will want to get information about travel habits, comfort levels and search patterns. An automobile manufacturer may need information about driving habits that can be obtained from on-board computers, the model type the consumer is interested in and brand awareness. This data may be derived from interactions with marketing campaigns, the use of tracking cookies, app activity, social network data or even surveys. The number of data fields can be enormous. The data are analysed, prepared and transformed into tangible information that supports marketing and communication.

Profiles help marketers to identify, understand and interact with customers on a granular, individual basis. It is possible to gain insights from the profile to drive decisions and create goal-targeted contextual, personalised interactions.

A profile, as we use it, is a single record (which may include thousands of variables) for an individual or an artificial representative of a small group of similar people. Such a profile is the result of a complex compiling process where data from different levels of individualisation and sources are brought together to give a good view of the individual or small group. A profile should be updated at regular intervals — with the timing based on how often changes happen. This individual information (raw data) must be stored in a big data solution or at least in a data warehouse. Ideally the raw data are kept for a certain time period, depending on the kind of information, so that it is possible to calculate the profiles for past periods as well. This allows learning from past campaigns and simulations of potential behaviour using earlier data to forecast planned campaigns.

Individual profiles allow marketing and sales to communicate with customers and prospects one-to-one. Every part of the campaign can be completely personalised to what the customer wants, to where they are in the buying cycle, and more. This, of course, requires an automated framework to compile a modular campaign or the installation of an alert system for trigger-based campaigns. This level of individualisation can be difficult and expensive, so most companies decide to use segments.

SEGMENTS

Segments are dynamic aggregations of homogeneous profiles. Among other things, cluster analysis and predictive models are used to calculate segments.

Segments help to operationalise communication and interaction, taking into account data from various available sources, such as social network data, existing master data management and CRM system data, behavioural data, transactional

data, etc. This makes personalised mass communication more flexible, without being overly complex. Segments provide an in-between strategy. One extreme is mass communication, where everyone in the selection gets the same communication. The other extreme is completely customised, individual communication, for example, based on individual profiles. Segments reflect the compromises. Several profiles can belong to a single segment or, if no specific profiles are built, all customers may be divided into several segments. So you might have 15 segments in a campaign, with segment-specific creative or tonality for each one.

It is especially important that the segments are actionable and customised, and always kept up to date.

IMPUTATION

Another advantage of segments and profiles is that they can help enhance and update customer data records using information found in the records of other profile members, but missing from a particular customer belonging to the same profile.

No database is perfectly accurate and we will always have bad data in our files, whether imputed or not. There is a risk involved in imputing data — but there is also a risk when we do not try to repair the data. The advantage of estimation based on other profiles or segments is that the risk of error is lower than estimating based on the full population, because the similarity between those records that have data and those without the specific data is provided in the profile. In this case, data from segments or profiles are used to generate an estimation of the missing value.

This is a common problem. As an example, birthday or age information is commonly missing. Experts distinguish between ‘real missing data’ and ‘not stored

information’. Typical examples of real missing data are missed gender and missed birthday or age information. Real missing data occurs in datasets generated from situations where the information itself exists in real life, but for unknown reasons is not stored in the dataset. We know that every customer actually does have an individual gender, and he or she was actually born on a particular date, even if those data are missing. The data may be missing because the customer explicitly decided not to share the information or because some process or steps to obtain the information did not work correctly.

If real missing information is detected, it can be replaced by estimation using imputation strategies. The methods differ depending on the business context, other existing information and the kind of information that is missed. Estimation can support replacement using the mean or median, or it may be done by a more complex method, such as a regression equation or a time series analysis.

An alternative way to replace missing values in real terms is to use third-party knowledge, matching the concrete dataset or the intelligent profile or segment the dataset belongs to.

The final result is an estimation, so there is a chance of being wrong. This might have an impact on the individual communication — you could miss the mark or, in extreme cases, even offend a potential customer. In general, we find that trying to complete our dataset is more advantageous than risky and delivers valuable knowledge for marketing.

HOW PROFILES AND SEGMENTS BECOME INTELLIGENT

Profiling and segmentation are not new — although the quantity of data we can use has grown exponentially, so our accuracy has improved. But these practices become

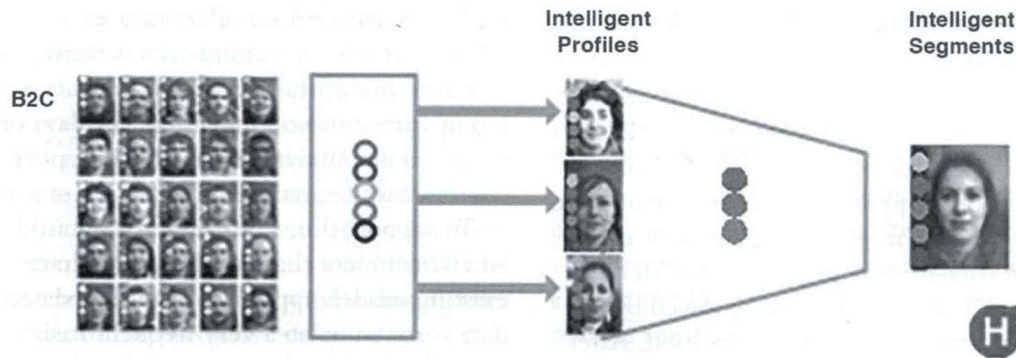


Figure 1: Intelligent profiles and segments in B2C

really exciting when they are combined with predictive analytics and other techniques as a basis for further analysis. Used in this way, they help to identify the potential for: cross-selling and upselling; better exhausting geographic opportunities; finding multipliers, influencers and brand ambassadors; optimising communication channels' mix and content; discovering trends early; reducing waste; counteracting churn risk and more.

Figure 1 illustrates how an intelligent profile represents a group of real people (with sparse datasets indicated by empty circles) and how several profiles will build a segment (with more complete data indicated by filled-in circles). If you decide that each customer will have his or her own profile, you will need layers. The decision ultimately depends on the type of business, the data collected and the number of customers in your database.

PREDICTIVE ANALYTICS: A SHORT INTRODUCTION

The heart of any kind of predictive analytics is the construction of predictive models.

Predictive modelling identifies patterns in the 'old data' in order to draw conclusions for future development. 'Old data' may include data from the previous day, from the last hour or even from the last click. In marketing, we particularly use information about the historical and current behaviour of customers and consumers to predict interests, needs and price acceptance, among other things.

Forecasting is always based on a known model using new but similar data, so the result of the deployment can provide a reasonable prediction about what is likely to happen (Figure 2).

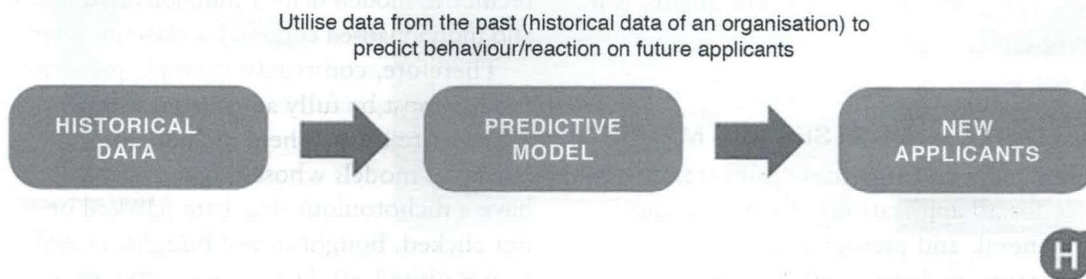


Figure 2: Learning from the past

KNOW WHERE TO GO

At the beginning of the analytical process there must be a precise formulation of objectives in order to select all the relevant data and get the best result. That is, the quality of the prediction depends mainly on asking the right questions. The next steps in the data-mining process deal with all the issues around the data, followed by the analytics and modelling phases itself. After a further step of validation and evaluation, a good model might be ready to be used.³

DATA PROCESSING AND FINE-TUNING

As soon as the objectives are defined, the available data are reviewed and interpreted. Referring to different time slots, new variables on different views are calculated per person. This results in a number of variables.²

The nature and scope of the data to be collected are defined in relation to the objectives. Often large volumes of data from diverse sources are only available 'raw', in a variety of formats, and are largely unstructured. Therefore, the data have to be processed, which means extracted, purified, filtered, analysed and transformed until they comply with the requirements, and can undergo the data-mining process for fine tuning. In this phase, data must also be tested for the need for more or different information, and whether the analytical concept must be optimised, among other things.

For adjustment and validation, test runs follow over the existing datasets. These are aggregated, adjusted and prepared. Only when the model is of sufficient quality is it fully used.

VELOCITY — EACH SECOND MATTERS

Time plays an immensely important role for all applications of data mining in general, and predictive modelling in particular: Incisive predictions must necessarily be present in almost real time

to deliver customised offers to users while they are on a company's website, at a booking portal or at a retail location. Companies can no longer wait for days or weeks to get analytical results to support current campaigns.

To support this, it is necessary to build an environment that enables you to have existing models applied as soon as updated data come in or on a very frequent basis. This also implies that the model itself allows fast application and that necessary data manipulation can be done in time, or that the data are already prepared so the new part can be added easily.

To design such an environment is complex and depends to a great degree on the actual usage of data, as well as the marketing and communication mix that is to be supported by predictive models. Some companies support this with management: They clarify where just-in-time/right-now prediction is relevant for business success and consumer satisfaction and where the marketing process does not justify just-in-time prediction. The result is a combination of approaches. One is a well-maintained, regular, batch process for data preparation and model application for all those models that support processes and campaigns where last-minute changes do not have any influence. And the other is a set of alert-based models, or predictive models calculated at the beginning of every web session. A typical alert-based example is credit risk prediction for new online or call-in customers. In a pure digital environment, targeted banners based on predictive models deliver individualised offers and individualised content for the consumer.

Therefore, constructing simple predictive models must be fully automated using appropriate tools. These include, for example, models whose target variables have a dichotomous structure (clicked or not clicked, bought or not bought, visited or not visited, student or not a student, etc). As part of predicting clicking or buying

behaviour, this is also used to impute missing values in datasets or profiles.

A certified subset of variables is used to obtain the stability, robustness and performance that are indispensable for automation. If a predictive model considers the last measured behaviour — for example, the current online session — it is essential that data can be accessed at any time during the session for the predictive calculation.

An essential step of profiling is data preparation. It is required for both modelling and verification, as well as for calculating score values. It is also the most time-critical process.

Automated predictive models cover a large proportion of tasks in banner optimisation by predictive targeting.³

A JOB FOR THE ANALYTICS TEAM

Complex predictive models and clusters are built by hand, in advance, by a team of analysts. This team is responsible for

constructing new variables, for converting unstructured data into analysis-capable datasets and for everything to which the complex rules of domain knowledge must be applied.

Records with many missing values or issues that require multi-class target variables mean that simpler analysis must be carried out by specialists in order to achieve reliable results.

Data mining methods such as descriptive statistics, cluster analysis, click-stream analysis, discriminant analysis, regression methods, decision trees, neural networks, case-based reasoning (CBR) and time series analysis are used.

Both automatically and manually built models are implemented into the existing architecture of a company so that they can take advantage of the existing environment and its synergies as much as possible.⁴

The environment shown in Figure 3 is able to develop, in a fully automated process, up to 1,500 new models a night.

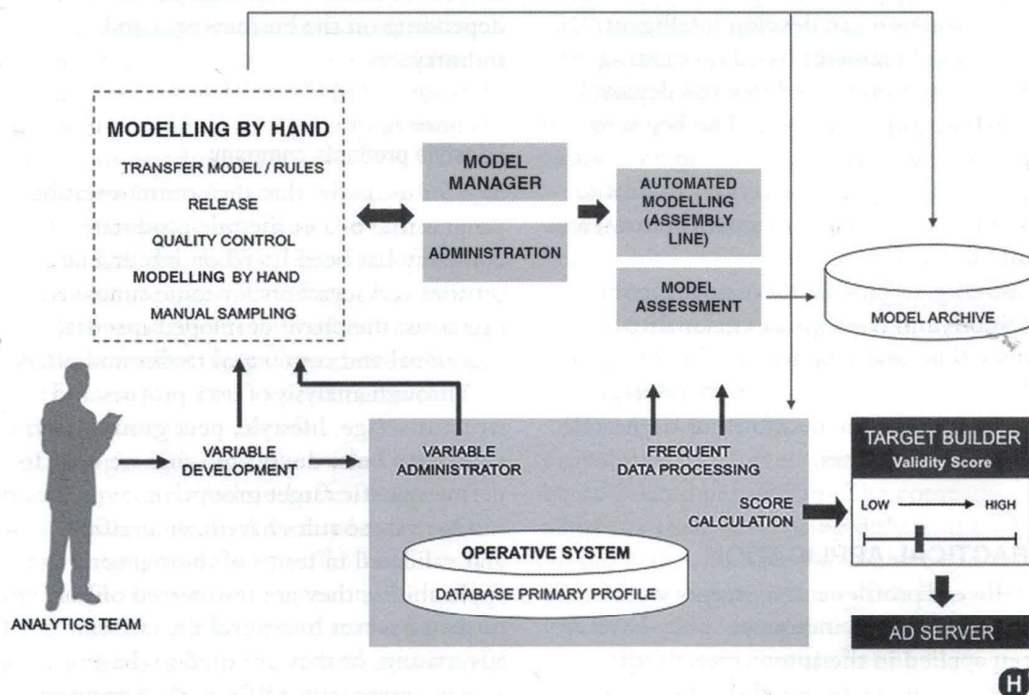


Figure 3: Modelling high-level architecture

The most complex part of the models is related to banners, but some processes are related to imputation (for example, estimating household income learned from online surveys to be estimated in profiles where those data are missing). All models are stored in the archive. Every active model goes through a regular assessment process to ensure quality. All active models are trigger based or applied just in time.

THE COMBINATION IS A CATALYST FOR THE RESULT

They are complex, time-consuming and labour-intensive, but the involvement of predictive analytics turns normal profiles and segments into intelligent ones. And these are tremendously valuable because they bring future development into play. Further analysis provides answers to company-specific issues that can make the difference between success and failure.

Even if it does not make sense under current business conditions to develop a sophisticated automated modelling system, an organisation can develop intelligent profiles and segments based on existing data preparation procedures and deploy models in batch processes. The key is to combine aggregated data descriptions with predictions of future behaviour to profile records (based on one or many persons) and build segments out of them.

Intelligent profiles drive goal-targeted decisions and contextual, personalised interactions and support. Intelligent segments support data-driven campaigns, promotional rationalisation and large-scale loyalty programmes.

PRACTICAL APPLICATION

Intelligent profiles and segments work in many different industries. They have been applied in the automotive sector (high propensity to invest) as well as in fast-moving consumer goods (FMCG)

companies. They are as indispensable in healthcare as they are for financial services providers.

If it makes sense to interact on a highly personalised scale with individuals or small groups of consumers, a framework is needed that will be automatically filled with campaign-relevant offers, text, creation and tonality based on each profile. We may end up with a couple of hundred different combinations to send out in the same campaign. Most of the time it is more relevant to individualise on the segment level and create an individual communication and follow-up programme by segment. Depending on the number of segments, you might have to design as few as, say, 15 different versions in one campaign.

While the use of intelligent profiles and segments will increase business, they are no replacement for testing. A variety of versions should be tested regularly for continuous improvement.

The principle remains the same: The objectives and the data needed vary depending on the business area and industry.

Lifestyle products company

Say, for example, that the communication programme of a lifestyle products company has been based on intelligent profiles and segments for some time. On this basis, they have developed specific, individual and contextual tactics and offers.

Through analysis of user profiles and structures (age, lifestyle, peer groups, etc), rules have been discovered that are used to define specific target groups.

Once these rules have been verified and validated in terms of their general applicability, they are transferred online to the ad server to control the current advertising, or they are used to drive e-mail campaigns. Offline, they provide the intelligence behind the selection

of addresses so that people will receive segmented, individual direct-mail communication.

Advertising material and banners are classified according to characteristics such as topic, format and content, and it is determined which potential target groups (segments) are to be served which advertising media in which time frame.

To determine the assignment of customers to intelligent segments, information about prototypical surfing behaviour, preferred communication channels and devices is used. We also include any concerns customers have revealed in previous purchases, requests, the social web or in personal contact.

If a person contacts the company for the first time, a profile is created with the first behaviour observed and data entered (for example, registration on the website with a postal address, landing page or surfing behaviour). If this consumer returns, the profile is supplemented with more data. In an online contact, the system detects the same person again and adds profile data from the current session. In an offline contact (call centre, letters, etc), relevant data are entered into the CRM system and will be used to map new persons as soon as possible to matching profiles.

All contact information is assigned to the appropriate profile in the CRM system. The threshold probability of a person belonging to one of the predefined profiles and segments is analysed regularly. This ensures that people can move among profiles and segments depending on their actual behaviour and needs.

During a campaign, the customer receives precisely tailored offers and advertising material. Any response, in turn, is expanding his or her profile. In this way, the profile is sharpened more and more from activity to activity and from response to response.

The communication programme retains the segment-based messages and is

steadily improved by adding profile-based communication, eg a price or rebate based on intelligent profiles.

The result: increased responses and sales

Customers and prospects receive better and better targeted communications that match their preferences. So the sales of our lifestyle product organisation increased by 3.5–4.0 per cent after the implementation of intelligent profiles and segments. And they continue to improve.

Automotive company example

What makes a prospective customer feel ready to buy a car? A 15 per cent discount? An extensive test drive? A free cup of coffee at the dealer? Would any of these have any greater effect than another?

To retain the best customers according to their customer lifetime value (CLV) and to attract new customers, customer loyalty programmes, sales incentives and offer packages must be truly tailored to the needs of website and point-of-sale visitors. The challenge is to increase customer satisfaction and maximise profits at the same time.

Thus, an automotive company decided to use intelligent profiles and segments about five years ago.

In addition to the standard profile data, other data were used: when the customer bought his last car, what model he owns, how often the vehicle was in for repair and why, which communication channels the customer prefers, what his hobbies are, what his calculated CLV is and more.

Based on these data, the purchase probability for a given (new) model can be calculated per person. The company addresses their advertising only to the 30 per cent who have the highest propensity to purchase or reach the next loyalty level. This means they save the cost of printing and distributing high-quality advertising brochures and avoid expenses for the dealers by reducing the




Figure 4: Intelligent profiles and segments in B2B

number of 'not truly serious' test drives. On the company's website, customers and prospective customers are shown the exact banners for which they have the highest affinity, according to their profiles. The space can be used more effectively and the potential customer feels welcomed and recognised.

The result: savings and increased sales

Since the application of intelligent profiles and segments, the automotive company has saved 28 per cent of the cost of product brochures and discounts. This is money they can spend on other marketing activities. The sales figures and return on investment (ROI) continue to rise steadily.

INTELLIGENT PROFILES AND SEGMENTS IN B2B

What's good for B2C customers and companies also supports B2B marketing.

Here, however, two levels need to be considered: the profiles of the companies (industry, products, size, turnover, structure, etc), as well as the profile of the relevant contact person(s) in the company (function, preferred communication channel, professional and personal interests, etc).

This is the reason why intelligent profiles and segments in the B2B sector are two-dimensional.

Depending on the objective, the company's and the contact persons' profiles are weighted differently.

If, for example, you want to stimulate interest in a new storage system, the company's profile is more important than that of the CEO, but the profile of the warehouse director might support the sales activity even more. When it comes to financial services, the financial situation and structure of a company surely matters, but it is also important to address the CEO individually.

If an advertising agency wants to bind customers or find new customers, they will devote maximum attention to the profiles of the marketing manager.

So using intelligent profiles and segments in B2B is complex, but if it is done right, it will speed up business outcomes (Figure 4).

Whether in B2C or B2B, intelligent profiles and segments must be kept alive and dynamic. That means they need to be cared for, updated and constantly fed the best information if they are to bring maximum benefit to consumers and businesses.

MINIMISE CRITICAL SUCCESS FACTORS

How do you feel when the shirt you ordered online a few months ago is offered to you today? Companies must adjust and optimise the data of the sales database and

the intelligent profiles and segments to avoid annoying customers and to really generate new business.

In other words, it is not enough to set up intelligent profiles and segments once, and then leave them to themselves.

There are some statistical challenges.

In selecting appropriate and purposeful forecasting methods, attention must be paid to the quality of the prediction, to stability and to the performance in the development. Concerning methods for quality assurance, success depends on robustness, run-time behaviour, parameterisation, fault detection and automation.

The requirements for transparency and data protection continue to increase. To collect, process and use data, you must take into account certain requirements at each stage of the process to conform to the respective national data protection laws. The further processing of data must ensure that the privacy of individuals is protected. This can be achieved by the use of sensitive technology.

CONCLUSIONS

The quantity and depth of accumulated and collected data have been discussed publicly for a long time. This means consumers' expectations of brands and companies, in terms of individualised, well-conceived and, in particular, really relevant communication, rise immensely.

Intelligent profiles and segments are used by more and more companies to provide

customers and prospects with what they (rightly) expect.

Companies will not only benefit from intelligent profiles and segments through the satisfaction and loyalty of their customers and prospects, but also through having found ways to monetise the knowledge gained by analysis. They may be able to operate lucrative affiliate marketing programmes, use insights for their own innovative management (eg data-driven products and services) or optimise processes and minimise fraud.

This added value will help companies using intelligent profiles and segments to remain competitive and sustainable. The better we get at creating profiles, the better the experience is for the consumer. One day we can have true one-to-one relevancy.

References

1. Li, L. and Memon, N. (2013) 'Mining groups of common interest: Discovering topical communities with network flows'. In: P. Perner (ed.). *Machine Learning and Data Mining in Pattern Recognition*. Springer-Verlag, Heidelberg, pp. 405–418.
2. Ahlemeyer-Stubbe, A. and Coleman, S. (2014) 'A Practical Guide to Data Mining for Business and Industry', John Wiley, Chichester.
3. Ahlemeyer-Stubbe, A. (2013) 'Informationen zur Optimierung des Predictive Behavioral Targeting'. In: M. Ceyp and J.-P. Scupin (eds). *Erfolgreiches Social Media Marketing*. Springer Fachmedien, Wiesbaden, pp. 181–182.
4. Ahlemeyer-Stubbe, A. (2010) 'Vollautomatisches Predictive Targeting und Modellierung des Reatime-Online-Verhaltens'. In: DDV (ed.). *Dialogmarketing Perspektiven 2010/11*, Gabler, Wiesbaden.