WORKING PAPER SERIES

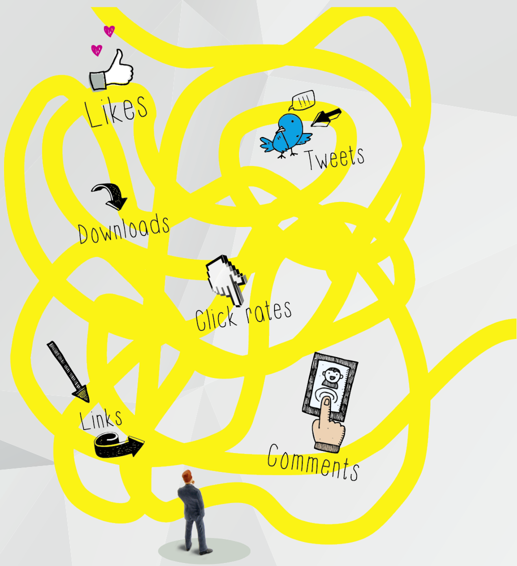

ANALYZING THE TEMPORAL DEVELOPMENT OF BRAND-RELATED SOCIAL MEDIA PHOTOS

A case study from the confectionary industry

Carolin Kaiser, Lara Enzingmueller,
Rene Schallner

WORKING PAPER /// NO. 6 / 2018

Analyzing the temporal development of brand-related social media photos

A case study from the confectionary industry

Carolin Kaiser*† Lara Enzingmueller‡ Rene Schallner†

Abstract— *A multitude of photos is posted each day on social media. These photos do not only reflect people's social lives but also their touchpoints with brands, representing a rich source of knowledge for market research. Since most of the early software tools focus on text, this wealth of information has long been neglected. Building on the new advancements in computer vision, this paper presents an automated approach for analyzing the temporal development of brand-related social media photos. The continuous analysis allows to detect consumer trends, to measure the influence of marketing campaigns and to examine the correlation with sales. Thus, marketers are in the position to identify opportunities and risks early on and to take appropriate actions. The approach is illustrated by a case study from the confectionary industry.*

Keywords— *Social media; User-generated photos; Image sharing; Visual eWOM; Computer vision; DCNN*

1 Introduction

The flood of photos on social media is overwhelming. Every day, 3.5 billion photos are exchanged via Snapchat[1]. 350 million photos are posted daily on Facebook[2] and 95 million photos per day are shared on Instagram[3]. The rapid increase in social media photos has been fostered by technological and social developments. Due to the ubiquitous availability of cameras and internet connections in mobile phones, people are able to take pictures at any time or place and share them with a large audience. At the same time, private photography is subject to change: Besides retaining memories and archiving special moments (Chalfen, 1987), people nowadays also capture everyday situations (Sarvas and Vihavainen, 2005). Sharing these snapshots with fellow users on social media has become an important mean of people's self-presentation and identity construction (Kapidzic, 2013).

These user-generated photos do not only reflect people's social lives but also their touch points with brands. For example, recently bought gadgets may be proudly presented to the camera, or favorite drinks may be part of photos from social gatherings. Thus, social media photos represent a valuable source of knowledge for market research which has not been fully exploited yet. Due to the multitude of photos, manual coding is not feasible, and automated solutions are crucial. Until recently, however, many software tools have focued on the analysis of textual postings only. New developments in the field of computer vision enable the automated analysis of social media photos. Especially Deep Convolutional Neural Networks (Krizhevsky et al., 2012) revolutionized the field of computer vision.

Considering the high relevance of social media photos and the recent advancements in computer vision, this work aims at leveraging the mostly unused potential of social media photos as a source of knowledge for marketing. The paper presents an approach that automatically gains marketing-relevant knowledge from user-generated photos. First, photos are collected from social media. Second, brand products and marketing activities are recognized in social media photos by applying Deep Convolutional Neural Networks. Third, the recognized products and marketing activities are analyzed with respect to their temporal development. The continuous analysis allows to detect important seasonal peaks and changing

*Corresponding author, carolin.kaiser@gfk-verein.org

†GfK Verein

‡University of Erlangen-Nuremberg

Citation: Carolin Kaiser, Lara Enzingmueller and Rene Schallner (2018), Analyzing the temporal development of brand-related social media photos—A case study from the confectionary industry, GfK Working Paper Series, No. 6 / 2018

[1]https://www.statista.com/statistics/257128/number-of-photo-messages-sent-by-snapchat-users-every-day/, accessed April 3, 2017

[2]https://www.omnicoreagency.com/facebook-statistics/, accessed April 3, 2017

[3]https://www.omnicoreagency.com/instagram-statistics/, accessed April 3, 2017

consumption trends, to measure the impact of marketing campaigns and to examine the coherence with sales. Thus, opportunities and risks for image management and sales can be identified at an early stage and corresponding actions in advertisement design and product development can be taken. The approach is illustrated by a case study from the confectionary industry and is based on approx. 50,000 photos for two confectionary brands posted on a popular image sharing platform over the course of two and a half years.

2 Related Work

Social media has changed the flow of information between companies and consumers. Companies have lost their communication monopoly. Users are no longer passive information seekers and recipients of companies' information (Tirunillai and Tellis, 2012). Through the ability to produce content, users have become active "prosumers" (producer + consumer), who share their experiences and points of view about products and brands (Hennig-Thurau et al., 2010). This brand-related user-generated content is perceived to be an especially trustworthy and credible source of information (Bickart and Schindler, 2001). Many studies have proven the influence of online reviews and posts on purchase decisions (e.g. Yayli and Bayram 2009, Xiaofen and Yiling 2009).

In addition to text-based content, ever more photos are uploaded to social media platforms as a form of socio-cultural expression (Huang and Park, 2013). Users enjoy sharing snapshots about their lives. Self-portraits, friends, activities and food are the most frequently posted forms of content (Hu et al., 2014). Photos vary with the users' age (White, 2010), gender (Smith and Cooley, 2012) and nationality (Huang and Park, 2013). By posting photos, users are able to express their personalities (Marwick, 2015). Besides information exchange and community support, status-seeking and self-representation have been found to be the main motives for photo sharing (Lee et al., 2015). Even if users seek to represent their ideal selves, social pressure from friends leads to a more realistic self-representation (Back et al., 2010).

Photos posted on social media provide not only deep insights into people's personal lives but also into their consumption preferences. Kaiser et al. (2016) show that photos posted to social media may be a sign of brand love. Papers from computer science demonstrate that it is possible to automatically derive user interests from people's posted photos (Xie et al., 2015), and to generate product recommendations (Feng and Qian, 2014).

Compared to textual posts, photos enable more intimate views on users' daily lives (Stefanone et al., 2010) and have a greater effect on the viewers. Studies from classic advertisement research have shown that photos are remembered longer (Shepard, 1967) and are more effective in changing a consumer's opinion about a brand than text (Mitchell and Olson, 1981). Furthermore, visual elements increase the perceived quality, credibility and authenticity of word of mouth messages. They reduce the insecurity about trusting product recommendations of other social media users (Lin et al., 2012). Shared photos are also more likely to catch people's attention and even motivate them to act (Lin and Huang, 2006). They generate greater product interest and purchase intention (Lin et al., 2012).

So far, the analysis of social media photos has mainly been addressed by research from psychology and computer science, but neglected by market research. Due to the increasing amount of posted social media photos, their rich information about consumer preferences and their high influence on fellow consumers, it is also important for market research to analyze them. To address this research gap (King et al., 2014), we propose an approach for monitoring brand-related user-generated photos and illustrate its benefits for marketing in a case study.

3 Approach

This paper proposes an approach for monitoring brand-related user-generated photos to gain marketing-relevant information. The approach consists of three main phases (see Figure 1). First, photos are collected from social media. Second, products and marketing-related objects are recognized in the collected photos by applying methods from computer vision. Finally, the detected objects are analyzed to gain valuable insights for marketing. The analysis focuses on the following three aspects:

- the temporal development of the posted products with respect to seasonality and trends
- the impact of marketing actions and external events on the temporal development on the posted photos
- the coherence between posted photos and sales

The continuous tracking of product images on social media enables an early detection of changes and risks as well as the implementation of appropriate marketing actions.

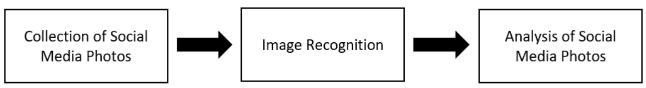

Figure 1: Approach

The approach is applied in a case study from the industry of fast-moving consumer goods. Due to short product life cycles, harsh competition, and high exchangeability of non-durable goods, up-to-date market research and quick reactions to market changes are especially crucial for companies of this industry. In this case study, we take two popular confectionary brands A and B as an example. To clearly define the market, we focus on German-speaking users. For this purpose, we collect photos from a popular image sharing platform which fulfill the following criteria:

- The brand name or a term highly associated with the brand must be mentioned in the hashtags or captions of the photos.
- German words must outweigh words of any other language in the hashtags, captions, comments and biography of the posting users.
- The photos posted between May 01, 2012 and October 31, 2015.

The final dataset comprises 49,074 photos: 35,073 photos for brand A and 14,001 photos for brand B.

4 Image Recognition

The aim of the image recognition is to detect the relevant products and advertisements in the photos by employing and extending state-of-the art computer vision methods. The detection is based on the R-CNN approach by Girshick et al. (2013) and comprises three main steps. First, potentially interesting region proposals, which might contain the relevant products or ads, are identified in the image. Second, features are extracted from each region proposal with the aid of Deep Convolutional Neural Networks (DCNN). Third, each region is classified whether it contains one of the relevant products or ads, based on the extracted features.

In order to find region proposals, the well-known algorithm *selective search* (Uijlings et al., 2013) is complemented with the VH-Connect algorithm developed by Winschel et al. (2016). Selective search exhaustively searches for good segmentations based on the structure of the image and has proven to be a good object independent algorithm for region proposals (Uijlings et al. 2013, Girshick et al. 2013). However, it has some problems to detect objects which do not have closed boundaries and consist of multiple components which is often the case for brand-related products. The VH-Connect algorithm is designed to generate additional proposals for such regions which are missed by selective search and has achieved good results on brand-related image sets (Winschel et al., 2016).

Features of the proposed regions are extracted by Deep Convolutional Neural Networks, which have recently become the most successful method in image recognition. Even though they have already been invented in the 1980s (LeCun et al., 1989), only the recent availability of large public image data sets and the increased performance capabilities of GPU computing enabled the successful transition from shallow feature encodings (Perronnin et al., 2010) to Deep Convolutional Neural Networks (Krizhevsky et al., 2012), derived from the visual cortex of animals (Hubel and Wiesel, 1968). We apply the VGG 16 network (Simonyan and Zisserman, 2014) which improves the original architecture of Krizhevsky et al. (2012) by pushing the depth of the network to 16 layers. By adding more convolutional layers, this network architecture won the localization and classification task of the ImageNet Challenge 2014.

Based on the extracted features of the region proposals, support vector machines are applied for classification. Support vector machines (Cortes and Vapnik, 1995) learn the parameters of a function that classifies the proposed regions in the best possible way. The goal is to identify the hyperplane, which separates the classes at maximum distance. In the simple case of two classes, the function can be depicted as a straight line seperating the two classes (product X versus not product X). Region proposals of the first class (product X) lie on one side of the line and region proposals of the second class (not product X) lie on the other side of the line. Support vector machines have been successfully applied in various classification tasks (Cai et al. 2004, Shin et al. 2005) for many years, and have achieved good results for the classification of image regions (Girshick et al., 2013).

The detection approach is based on supervised learning. This means, that before products and marketing-related objects can be detected, the algorithm must first be trained first with the aid of manually annotated training images. Since DCNNs consist of a large number of parameters, huge data sets are required for training. To reduce the manual effort of annotating a multitude of training images, we use an existing DCNN trained on the large auxiliary dataset ImageNet (Deng et al., 2009), and fine-tune it to our domain, which has proven to be an effective transfer learning paradigm for DCNN learning (Girshick et al., 2013).

To train and test the detection of products and advertisements, 170 photos per class were collected and all depicted products and marketing-related objects were manually encircled by a human annotator. For each class, 150 images were used for fine-tuning the DCNN, as well as training the SVMs, and 20 images were used to test the detection performance of the system. The performance test yielded a very good average precision of 0.95 and recall of 0.76. This means that 95% of the detected products and marketing-related objects are correct and 76% of all products and marketing-related objects are found. Detailed results per product class are shown in Table 1.

Brand	**Object Classes**	**Precision**	**Recall**
Brand A	Chocolate Hearts	1.00	0.80
	Chocolate Eggs	1.00	0.88
	Advent Calendar	0.88	0.71
	Chocolate Bar	0.96	0.71
	Chocolate Santa	1.00	0.87
	Chocolate Bunny	1.00	0.83
	Brand Mascot	0.89	0.64
Brand B	Advent Calendar	1.00	0.83
	Chocolate Cubes	0.94	0.67
	Chocolate Bar	0.76	0.79
	Mini Chocolate Bars	0.96	0.62
	Ad Poster	0.95	0.75
	Ad at Unique Place	1.00	0.80
	Ad in Unique Form	1.00	0.77
	MEAN	**0.95**	**0.76**

Table 1: Performance

5 Image Analysis

The detected products and marketing-related objects are aggregated by time and analyzed with respect to the temporal development. For companies, these analyses can deliver valuable information: Trends and peaks can be detected, the success of marketing actions can be measured, and chances and risks for brand image and brand sales can be estimated at an early stage.

5.1 Seasonality and trends

At first, the temporal development of brand related social media photos is examined with respect to seasons and trends. The continuous tracking enables companies to measure their visual buzz and to understand consumers' posting behavior. Thus, marketers are in a better position to plan their marketing strategy or to adjust their product assortment.

The temporal development of all posted pictures for brand A and brand B shows that the number of posted photos increases for both brands (see Figure 2). This can be traced back to the growing user base of image sharing platforms and the spread of mobile phone usage.

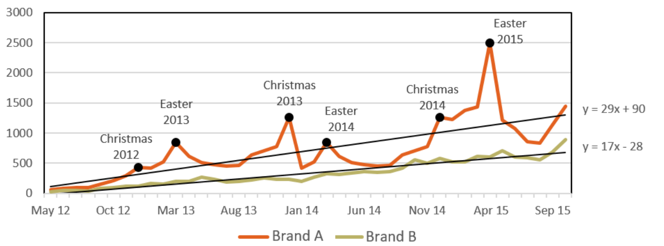

Figure 2: Temporal development of the number of posted photos for brand A und B

Even if the number of photos increases for both brands, several differences can be noticed between them: The trend line of brand A increases more strongly than the one for brand B, indicating a faster growing and a more engaged user community for brand A. Furthermore, the time series of brand A shows seasonal peaks while the time series of brand B is almost free from fluctuation. This means that the posting activities for brand A vary to a much higher degree than the ones for brand B. This difference might be attributed to the different product ranges of both brands. While brand B only offers some seasonal varieties of chocolate bars, the product range of brand A includes a large variety of season specific chocolate products which triggers users to take and share photos, especially at annual seasonal events. One significant periodic peak in the time series is Easter Sunday (see Figure 3). The reason for this peak is the high number of posted Easter-related products. The number of shared Easter bunnies and chocolate Easter eggs reaches its climax on Easter Sunday of every year. Users enjoy taking pictures of their Easter presents to show what they got and to send out their individual Easter greetings to others. Therefore, the peak is day-specific. The number of posted Easter products on this special day amounts to a multiple of the average of the other Easter holidays. After reaching the climax, the number of posted products decreases again.

Further peaks can be found in December during Christmas season (see Figure 3). A high number of pictures of brand A are posted on the 6th of December in all considered years. While the overall quantity of pictures showing Santa Clauses increases every year, the relative share of them is almost constant, ranging from 38% to 42%. Besides, the heart-shaped chocolate from brand A is another product for which the time series exhibits day related peaks at Valentine 's Day and Mother's Day.

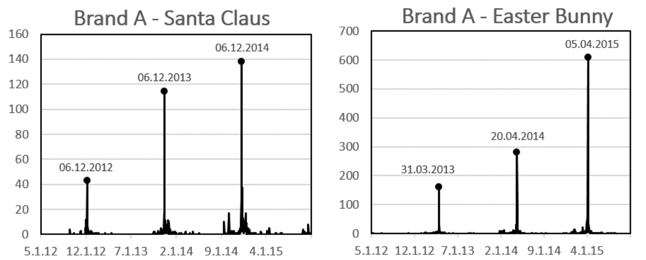

Figure 3: Temporal development of posted photos showing seasonal products from brand A

In contrast to brand A, which offers a broad variety of seasonal products, brand B focuses on traditional chocolate bars. This strategy is also reflected by people's posting behavior. While 43.7% of all posted pictures about brand B include at least one bar of chocolate, the relative share of chocolate bars of all pictures about brand A is only 25.9%. However, brand B does not totally refrain from offering seasonal products. Their product range includes seasonal bars in form of special summer-editions and winter-editions. Analyzing the meta data of the pictures revealed that the summer-editions are mainly posted in summer time, whereas the winter editions are mainly posted in the colder time of the year. This means, that with every new season consumers are motivated again to take and post photos of the seasonal chocolate bars. Especially in the last year of the relevant period, the peaks are significantly high. The rising number of social influencers posting special varieties might be the reason for the significant increase.

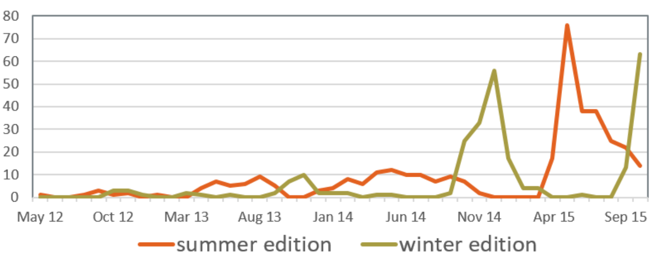

Figure 4: Temporal development of posted photos showing seasonal chocolate bars from brand B

As Figure 4 shows, brand B manages to keep the posting activity of seasonal products on a higher level for a longer period while seasonal products of brand A result in day-specific peaks. With its seasonal chocolate bars, brand B receives less visual buzz on special holidays like Christmas and Easter but accomplishes to keep the social media activity on a high level, even during the summer months which are usually weak months for confectionary brands.

Analyzing the meta data for the pictures of brand B also reveals changing user preferences. For several years, brand B produced some chocolate bars which did not contain any animal products, but did not promote this fact. However, the keyword 'vegan' was increasingly mentioned by the social media users between 2012 and 2015. In 2016, brand B reacted to the users' new food preferences by developing two special vegan varieties, especially promoting that they are vegan.

5.2 Impact of marketing actions and external events

Analyzing the temporal development of posted pictures does not only reveal seasons and trends but also delivers valuable information about the impact of marketing actions and external events on users' posting behavior. Thus, the effect of marketing campaigns can be evaluated at an early stage and adapted quickly if needed. The impact is measured in two ways. The direct measure considers whether pictures of the marketing actions are shared on social media. The indirect measure evaluates if these actions trigger social media users to post photos of all kind of brand-related products.

Brand B mainly uses uses out-of-home advertising. During spring and autumn, huge advertising posters are affixed at central locations in several German cities. These posters show the products together with creative and humorous slogans. The analysis reveals that the amount of advertisement pictures is especially high during the months March-June and September-November when the posters are displayed (see Figure 5). The humorous posters seem to motivate people to take pictures and share them immediately on social media.

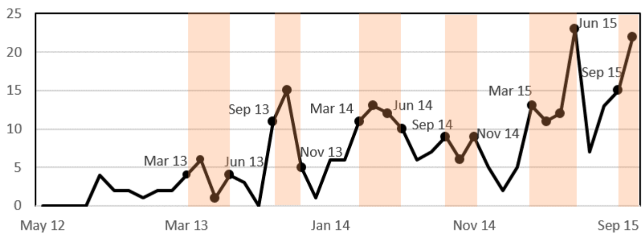

Figure 5: Temporal development of posted photos showing ad posters of brand B

Even though brand A has also displayed advertisements on posters between 2012 and 2015, the users of the image sharing platform do not post them. Since summer 2016, brand A has started placing their advertisment posters on similar locations as brand B. This might have been a reaction to the successful branding of brand B and other brands.

Besides poster advertisements, brand B uses uses unique places and forms for their out-of-home advertising which have evolved into an attraction people want to pose with for a photo and share it online. An increasing number of user-generated social media photos contain these unique types of advertisement.

For a long time, brand A has established a mascot which is well recognized by German consumers. Nowadays, brand A organizes special events with the mascot in German cities and asks people to share pictures of the event. A lot of social media related events took part in the considered period and led to an increase in pictures showing the mascot. Brand A uses the popularity of the mascot also for point-of-sale marketing. A lot of pictures show people hugging or kissing the mascot, which demonstrates the existing emotional connection between people and this mascot and eventually the brand itself. Brand A also positions the mascot at the entrance of their own stores. When the one of their major stores closed, less pictures of the mascot were posted on social media. Compared to the month before, the number of generated pictures showing the mascot decreased by 41%.

Both brands succeed in motivating people to take and share photos of their marketing actions. However, only a small number of marketing events lead to an increase in the overall number of brand related pictures. When looking at the number of posted photos of brand B, three major peaks can be attributed to social media campaigns inviting social media users to visit special events or to test new products, whereas two main peaks occur at seasonal events (see Figure 6). Even without offering seasonal products, brand B benefits from posted pictures showing seasonal products of other confectionary brands next to their own products. Within the considered period, the brand even organized an event, which shows the company's clear dissociation of seasonal products: they placed a machine at a train station, which changes Santa Claus shaped chocolate into chocolate bars.

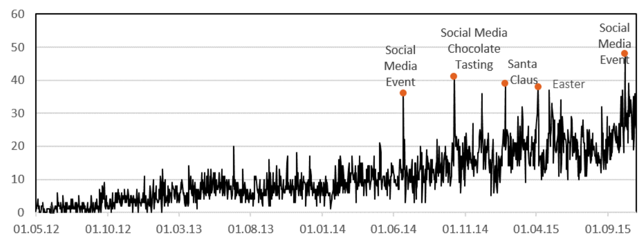

Figure 6: External events and peaks of posted photos for brand B

In contrast to brand B, the time series of brand A is strongly shaped by seasonal events (see Figure 7). Peaks are reached at Easter Sunday, Saint Nicholas Day, Valentine's day, and the beginning of the advent season. There are no major peaks at social media related campaigns or special events such as the sponsoring of sports championships or the celebration of the brand's anniversary.

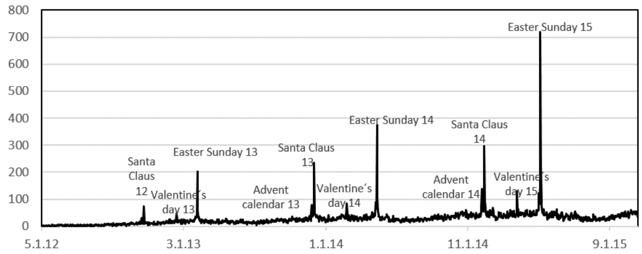

Figure 7: External events and peaks of posted photos for brand A

5.3 Coherence with Sales

For companies it is important to know what impact posted pictures have on their success in the future. Particularly the question, how user-generated pictures of a company influence sales is essential. For this purpose, we compare the absolute number of brand specific postings with sales figures. The sales data of both brands is sourced from a panel with 30,000 households of a leading market research company in Germany.

The analysis is performed on a weekly basis. Sales is measured as the number of households that purchased products of the brands within a specific week. Figure 8 and Figure 9 illustrate the temporal development of sales and posted pictures. Sales is only depicted in form of an index for data protection purposes. For each calendar week, the number of households which bought products of the brands are divided by the number of households which bought products of the brands in the first week of the period under consideration. The number of pictures for brand B is given in per cent, and for brand A in per mill to achieve a clear presentation.

Before analyzing the correlation between sales and social media postings, it is necessary to examine the time series for trends and seasonality to ensure that the measured correlation does not represent a spurious relationship caused by the coincidence of a third unseen factor (Greene, 2011). The time series for posted pictures and sales of brand B show an upward trend (see Figure 8). During the considered period brand B seems to get more popular. More people are buying the products of brand B and more users are posting photos about brand B.

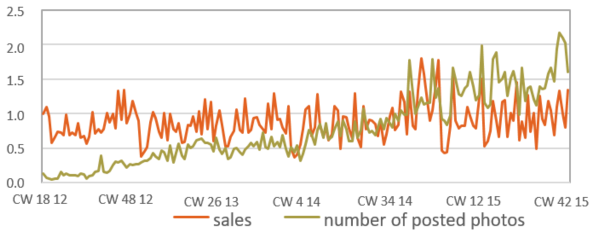

Figure 8: Temporal development of sales and number of posted photos for brand B

Looking at the graphs of brand A, it is apparent that sales and posted pictures both follow a seasonal structure with peaks at Easter and Christmas (see Figure 9). A huge number of seasonal products of brand A are bought and photographed during the holiday season. The difference between the two graphs of brand A is, that sales start increasing earlier before the peak than the number of posted photos. The upward trend of brand B and the seasonality of brand A are eliminated from the time series before applying a cross-correlation-analysis.

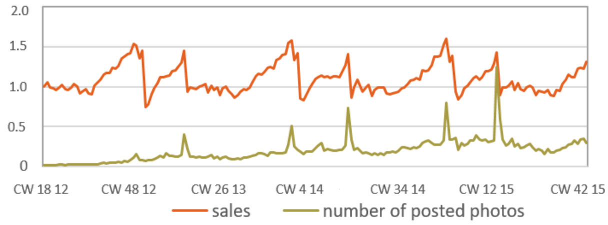

Figure 9: Temporal development of sales and number of posted photos for brand A

Cross-correlation-analysis determines whether a time series can be used as a prognostic indicator for another time series. The cross-correlation-coefficient (CCF) measures the degree to which two time series correlate at a certain time-lag. The analysis reveals a significant correlation between sales and posted pictures for both brands (see Figure 10). On a 99% confidence level, those two variables correlate at lag 0. This means, that an increasing number of posted pictures in one week goes along with increasing sales in the same week. However, it is not possible to make a statistical statement whether user-generated postings encourage other users to purchase the brand's products, or if the pictures are posted after the products have been bought.

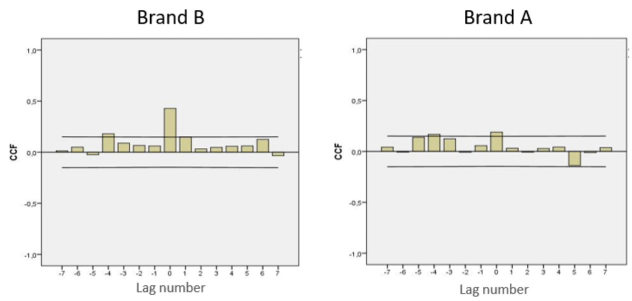

Figure 10: Cross-correlation between sales and the number of posted pictures

It can be assumed, that both directions of temporal correlation exist in our case study. Consumers purchase confectionary goods and post pictures of the acquired products. These brand-related pictures impact the purchase decisions of other users and might influence them to buy confectionary products. User-generated photos are a kind of visual eWOM (Ahuvia et al. 2017, Kaiser et al. 2016) and a reliable source of information for fellow users (Lin et al., 2012). Since these photos often contain the buyer's honest opinion concerning the depicted product, they have the potential to influence purchase decisions (Lin et al. 2012, Lin and Huang 2006). The analysis of the pictures' meta data shows, that in this case study, both directions of temporal correlation exist. Users of the image sharing platform purchase products and show them to their online community like this example: *"Bought this at Edeka today. They are sooo delicious. Highly addictive! [...] "*. Other users mention that they were motivated to buy a certain product by seeing it: *"You made me do it! Yes, I admit it, you disposed me to do it... often walked past and now I finally bought it! I am curious about it. [...] "*

Further knowledge about the direction of the temporal correlation between sales and posted pictures might be gained by analyzing the data on a daily basis. However, the dataset at hand was not large enough to conduct such an analysis.

6 Conclusion

The number of pictures posted on social media is increasing steadily. These pictures do not only reflect people's social lives but also their experiences with brands, thus representing a rich source of knowledge for companies. By continuously monitoring the content of the posted pictures, consumer trends can be identified, the impact of marketing actions can be measured and the coherence with sales can be estimated. However, due to the huge number of pictures, a manual analysis is not feasible. This paper presents an approach which enables the automated analysis of pictures by computer vision and highlights the economic relevance with a case study in the confectionary industry.

The analysis of brand-related products in this case study shows, that even if both brands are pursuing a different strategy concerning their product range, both are operating successfully. While seasonal products of brand A are massively bought and shared in the holiday season, brand B is focused on chocolate bars and achieves to keep sales and posted pictures almost stable during the year. Besides posting photos of actual products, the users of the image sharing platform also post advertisement elements of the two considered brands. Creatively and emotionally designed advertisements, as well as social media related marketing actions, result in a higher number of posted pictures.

The short-term effects of visual content on performance indicators is difficult to quantify. The quantitative analysis shows, that pictures and sales are dependent on each other, while the temporal direction of correlation cannot be determined. The qualitative analysis of the meta data indicates that both temporal directions of correlation can be found in the case study. Social media users are posting pictures of chocolate products they have bought which may then trigger fellow users to also buy these products. As a form of electronic word of mouth, brand-related social media photos reflect the users' opinions about specific products and can influence the purchase decisions of other users.

Because German chocolate consumption has been constant for years and the market is strongly competitive, confectionary manufacturers must choose a clear strategy to be able to succeed in the future. The temporal analysis of social media pictures enables marketers to detect new consumption trends, to evaluate marketing campaigns, and to estimate purchase activities. Thus, chances and risks can be identified at an early stage and corresponding actions, such as the organization of social media events or the integration of influential consumers, can be planned. The present interaction of posted pictures and sales highlights the high potential of social media for marketing.

Future research is crucial to shed light on the causal-effect-chain between picture posting and sales. While this study was based on posting data and sales data from two different sources, further insights can be expected from a single source approach, where information about posting and purchasing is collected from the same group of people. Further investigations could also determine to what extent the numerous postings of innovative products influence the sales of classic ones. Besides pictures, the analysis of other forms of user-generated content has great potential for further research. Especially short video clips posted on social media platforms enjoy great popularity and represent a valuable source of consumer information for market research.

Acknowledgement

The authors would like to thank Rainer Lienhart, Christian Eggert, and Philipp Harzig from the University of Augsburg for the excellent cooperation in image recognition.

References

Ahuvia, A., Kaiser, C., Rauschnable, P., and Wimble, M. (2017). Corporate Branding and Facebook Brand Pictures. In *Proceedings of the 3rd International Colloquium on Corporate Branding, Identity, Image and Reputation*, London.

Back, M. D., Stopfer, J. M., Vazire, S., Gaddis, S., Schmukle, S. C., Egloff, B., and Gosling, S. D. (2010). Facebook Profiles Reflect Actual Personality, Not Self-Idealization. *Psychological Science*, 21(3):372–374.

Bickart, B. and Schindler, R. M. (2001). Internet forums as influential sources of consumer information. *Journal of Interactive Marketing*, 15(3):31–40.

Cai, Y.-D., Ricardo, P.-W., Jen, C.-H., and Chou, K.-C. (2004). Application of SVM to predict membrane protein types. *Journal of Theoretical Biology*, 226(4):373–376.

Chalfen, R. (1987). *Snapshot Versions of Life*. University of Wisconsin Press.

Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3):273–297.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.

Feng, H. and Qian, X. (2014). Mining user-contributed photos for personalized product recommendation. *Neurocomputing*, 129:409–420.

Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2013). Rich feature hierarchies for accurate object detection and semantic segmentation. *arXiv:1311.2524 [cs]*.

Greene, W. H. (2011). *Econometric Analysis*. Pearson Higher Education, Boston, Mass., 7th revised edition edition.

Hennig-Thurau, T., Malthouse, E. C., Friege, C., Gensler, S., Lobschat, L., Rangaswamy, A., and Skiera, B. (2010). The Impact of New Media on Customer Relationships. *Journal of Service Research*, 13(3):311–330.

Hu, Y., Manikonda, L., and Kambhampati, S. (2014). What we instagram: A first analysis of instagram photo content and user types. In *Proceedings of the 8th International Conference on Weblogs and Social Media, ICWSM 2014*. The AAAI Press.

Huang, C.-M. and Park, D. (2013). Cultural influences on Facebook photographs. *International Journal of Psychology: Journal International De Psychologie*, 48(3):334–343.

Hubel, D. H. and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, 195(1):215–243.

Kaiser, C., Ahuvia, A., Rauschnabel, P., and Wimble, M. (2016). Visual eWOM: Are Facebook Brand Photos a Sign of Brand Love?

Kapidzic, S. (2013). Narcissism as a predictor of motivations behind Facebook profile picture selection. *Cyberpsychology, Behavior and Social Networking*, 16(1):14–19.

King, R. A., Racherla, P., and Bush, V. D. (2014). What We Know and Don't Know About Online Word-of-Mouth: A Review and Synthesis of the Literature. *Journal of Interactive Marketing*, 28(3):167–183.

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, NIPS'12, pages 1097–1105, USA. Curran Associates Inc.

LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., and Jackel, L. D. (1989). Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Computation*, 1(4):541–551.

Lee, C. S., Bakar, N. A. B. A., Dahri, R. B. M., and Sin, S.-C. J. (2015). Instagram This! Sharing Photos on Instagram. In *Digital Libraries: Providing Quality Information*, Lecture Notes in Computer Science, pages 132–141. Springer, Cham.

Lin, T. M., Lu, K.-Y., and Wu, J.-J. (2012). The effects of visual information in eWOM communication. *Journal of Research in Interactive Marketing*.

Lin, Y.-S. and Huang, J.-Y. (2006). Internet blogs as a tourism marketing medium: A case study. *Journal of Business Research*, 59(10):1201–1205.

Marwick, A. E. (2015). Instafame: Luxury Selfies in the Attention Economy. *Public Culture*, 27(1 (75)):137–160.

Mitchell, A. A. and Olson, J. C. (1981). Are Product Attribute Beliefs the Only Mediator of Advertising Effects on Brand Attitude? *Journal of Marketing Research*, 18(3):318–332.

Perronnin, F., Sánchez, J., and Mensink, T. (2010). Improving the Fisher Kernel for Large-Scale Image Classification. In *Computer Vision – ECCV 2010*, Lecture Notes in Computer Science, pages 143–156. Springer, Berlin, Heidelberg.

Sarvas, R. and Vihavainen, S. (2005). A Holistic View on Future Snapshot Media.

Shepard, R. N. (1967). Recognition memory for words, sentences, and pictures. *Journal of Verbal Learning and Verbal Behavior*, 6(1):156–163.

Shin, K.-S., Lee, T. S., and Kim, H.-j. (2005). An application of support vector machines in bankruptcy prediction model. *Expert Systems with Applications*, 28(1):127–135.

Simonyan, K. and Zisserman, A. (2014). Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv:1409.1556 [cs]*.

Smith, L. R. and Cooley, S. C. (2012). International Faces: An Analysis of Self-Inflicted Face-ism in Online Profile Pictures. *Journal of Intercultural Communication Research*, 41(3):279–296.

Stefanone, M. A., Lackaff, D., and Rosen, D. (2010). Contingencies of Self-Worth and Social-Networking-Site Behavior. *Cyberpsychology, Behavior, and Social Networking*, 14(1-2):41–49.

Tirunillai, S. and Tellis, G. J. (2012). Does Chatter Really Matter? Dynamics of User-Generated Content and Stock Performance. *Marketing Science*, 31(2):198–215.

Uijlings, J. R. R., van de Sande, K. E. A., Gevers, T., and Smeulders, A. W. M. (2013). Selective Search for Object Recognition. *International Journal of Computer Vision*, 104(2):154–171.

White, L. (2010). Facebook, Friends and Photos: A Snapshot into Social Networking for Generating Travel Ideas. *Tourism Informatics: Visual Travel Recommender Systems, Social Communities, and User Interface Design*, pages 115–129.

Winschel, A., Lienhart, R., and Eggert, C. (2016). Diversity in Object Proposals. *arXiv:1603.04308 [cs]*.

Xiaofen, J. and Yiling, Z. (2009). The Impacts of Online Word-of-mouth on Consumer's Buying Intention on Apparel: An Empirical Study. In *International Symposium on Web Information Systems and Applications*.

Xie, P., Pei, Y., Xie, Y., and Xing, E. (2015). Mining User Interests from Personal Photos. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, pages 1896–1902, Austin, Texas. AAAI Press.

Yayli, A. and Bayram, M. (2009). eWOM: The Effects of online consumer reviews on purchase decision of electronic goods.