

Harnessing AI to Navigate User Generated Content: A Framework for Consistent Customer Experience Across Digital Touchpoints

Natalie Rubin

Research Associate

Nuremberg Institute of Technology, Faculty of Business Administration

90489 Nuremberg, Germany

natalie.rubin@th-nuernberg.de

Philipp Haan

Research Associate

Fraunhofer Institute for Integrated Circuits IIS

90489 Nuremberg, Germany

philipp.haan@iis.fraunhofer.de

Ralph Blum

Professor for General Business Administration

Nuremberg Institute of Technology, Faculty of Business Administration

90489 Nuremberg, Germany

ralph.blum@th-nuernberg.de

ABSTRACT

The dynamic nature of the ever-increasing amount of user-generated online content (UGC) creates a host of new digital touchpoints for contacts between companies, consumers, and other stakeholders. While some of these touchpoints can be directly influenced by enterprises, the control and content design of so-called earned touchpoints, like product reviews or brand related social media posts, lies outside their sphere of influence. Aligning the content of these touchpoints is crucial for creating a unified brand image and enhancing brand success. This study aims to develop and test a methodological framework that recognizes the uncontrollability of external touchpoints from an enterprise perspective and supports marketing practitioners in delivering a consistent consumer experience (CE) across all touchpoints. The methodology employs AI-powered techniques for automated topic extraction using natural language processing (NLP) and image recognition (IR) to analyze consumer generated content. These techniques form the basis for informed decision-making and resource allocation in both strategic and operational aspects of customer experience management (CEM). The findings indicate that the automated approach to UGC-analysis is suitable for enhancing the understanding of the CE and guiding the design of brand-owned touchpoints.

Keywords: Consumer Insights, Customer Journey, Natural Language Processing, Image Recognition, Digital Touchpoints

1 Introduction

In recent decades, digitization has significantly reshaped how consumers communicate with and about companies and their products, and which touchpoints they encounter along their customer journey (CJ). This increased connectivity with consumers creates an opportunity for companies to identify customer needs and to design personalized products and services by utilizing this knowledge, radically transforming the customer experience (CE) (Hoyer *et al.*, 2020). At the same time, enterprises are losing their monopoly over the creation and communication of brand-related information, as discussions shift to user-centric platforms and social media, enabling peer-to-peer consumer conversations (Liu *et al.*, 2020). As a requirement for effective marketing design, “[...] researchers must be able to extract underlying insight - to measure, track, understand, and interpret the causes and consequences of marketplace behavior” (Berger *et al.*, 2020). Finding relevant user-generated content (UGC) within the vast amount of dynamically created data and applying efficient analysis methods for fast insight generation poses a challenge for marketing practitioners.

In response, this paper proposes and tests a structured, semi-automated methodology that focuses on eliciting brand-related content from UGC, originating from non-brand-owned pre-purchase touchpoints, with the aim of creating a consistent and thematically cohesive CE. The methodology utilizes natural language processing (NLP) and image recognition (IR) to identify distinct and semantically cohesive topics from UGC, focusing primarily on the consideration stage of the consumer journey. While existing studies demonstrate significant interest in automated text and image analysis in marketing, the scientific focus, to the best of our knowledge, has not yet been on the CE framework.

The effectiveness of our method is tested using UGC (product reviews and brand related image-posts) dealing with a brand of the handicraft domain, namely the air-drying clay “FIMOair” produced by Staedtler SE¹. Our approach successfully identified approximately 70 percent of the discussed topics, helping to understand and aggregate the brand related communication along the CE framework.

2 Theoretical Background on CJ, CE, CEM and UGC

To optimize the allocation of (online) marketing resources in accordance with the relative consideration and purchase impact of different channels, a profound understanding of the CJ is vital (Wolny & Charoensuksai, 2014). Within the CJ, the path of purchasing behavior contains a multitude of touchpoints, which can be summarized as the dynamic and multidimensional construct of CE (Lemon & Verhoef, 2016, 2016). This paper focuses on supporting customer experience management (CEM), which involves effectively composing these touchpoints and coordinating their design and deployment, alongside the creation of new touchpoints (Hoyer *et al.*, 2020). Despite the increasing complexity of CJ, the thematic cohesiveness and consistency of the various interaction platforms plays a key role in brand success determinants, such as customer loyalty (Homburg *et al.*, 2017; Kuehn *et al.*, 2019; Keyser *et al.*, 2020). In this context, thematic cohesion refers to “[...] the extent to which consumers perceive multiple touchpoints as sharing a common brand theme” (Kuehn *et al.*, 2019).

¹ <https://www.staedtler.com/uk/en/products/fimo-modelling-clay-accessories/fimoair/>

Although a typical CJ consists out of three stages: pre-purchase, purchase and post-purchase (Lemon & Verhoef, 2016), our focus is primarily on the pre-purchase stage, as it has the most significant impact on the consumers buying decision (Wang *et al.*, 2017). This stage is determined by the extent to which consumers contemplate a potential purchase of the brand's products, known as "consideration" (Keyser *et al.*, 2020).

Generally, CEM must consider three types of touchpoints in the context of brand-consumer interaction: paid (partner-owned), owned (firm-owned) and earned (social) media (Stephen & Galak, 2012). While brand-owned touchpoints refer to platforms, such as firm websites or loyalty programs, whose content can be fully controlled by the enterprise (Lemon & Verhoef, 2016; Nam & Kannan, 2020), paid or partner-owned touchpoints represent channels that are jointly designed by a company and additional partners, such as multichannel distributors (Stephen & Galak, 2012; Lemon & Verhoef, 2016). In contrast, social or earned touchpoints, originating from brand-related consumer-to-consumer communication and interaction, cannot be controlled or influenced by a company and instead rely on UGC shared with the community. Examples of earned media include social networks and online reviews (Stephen & Galak, 2012).

As consumer experiences become more diverse and complex, peer-to-peer interactions gain momentum, influencing the usage and overall effectiveness of other touchpoints throughout the CJ and shaping online community activity (Stephen & Galak, 2012; Barwitz & Maas, 2018). These social touchpoints are of particularly high importance in the pre-purchase and purchase stages of the CJ, when consumers are in the process of elaborating their purchase-decision, seeking guidance (Nam & Kannan, 2020). Baxendale *et al.* (2015) conceptualize two main types of peer-to-peer or social touchpoints. Electronic word-of-mouth (eWOM) can be understood as "[...] consumer-generated, consumption-related communication that employs digital tools and is directed primarily to other consumers" (Babić Rosario *et al.*, 2020). Furthermore, consumers' brand attitudes can be influenced by others through mere observation, a phenomenon known as "peer-observation" (Baxendale *et al.*, 2015). As communication shifts to the online environment, these observations become centered around social media, where users present their purchases and use-case scenarios.

Acknowledging the limitations of earned channels due to their lack of controllability, research is needed to unveil how a consistent CE can be established, enabling a thematically cohesive interconnection of earned, paid, and owned channels. Therefore, the methodology introduced in this paper investigates insight extraction from UGC, such as eWOM and peer-observation, with the goal of harmonizing the information with company-controlled owned and paid channel content.

3 Related Work on Automated UGC Analysis in Marketing

User-generated Content (UGC) is defined as publicly accessible consumer-created data that relies on the creators' creativity and is generated without monetary incentives or professional practices (Tirunillai & Tellis, 2012). A key characteristic of UGC is reflected in its unstructured nature, resulting in non-numeric and multifaceted information providing a range of unique insights and thus simultaneously representing several phenomena (Balducci & Marinova, 2018). Investigating UGC, specifically in the form of brand-related online communication, holds high potential to support marketing decision-makers. The abundance of unaided and presumably unbiased data underpins its scalability, while its

quick and comparatively inexpensive availability ensures cost-efficiency (Timoshenko & Hauser, 2019; Dzyabura & Peres, 2021). Recent studies unveil how UGC can be employed in the field of marketing, serving as at least equally valuable source for the elicitation of customer insights compared to traditional survey methods (Timoshenko & Hauser, 2019). Similarly, brand-related online communication, such as UGC, significantly contributes to a positive brand perception, crucial for the development of consumer-based brand equity (CBBE) (Keller, 2013). Furthermore, state-of-the-art research highlights UGC as a driver for brand awareness and satisfaction (Colicev *et al.*, 2019). However, the challenge lies in the large quantity and variety of unstructured data shared by customers and stakeholders, necessitating data-driven and automated solutions for effective analysis of UGC.

The automated analysis of brand- and product-related online content has been a topic of interest in the marketing domain throughout the recent years (Berger *et al.*, 2020). For insight generation, automated topic detection can serve as a tool that helps to unveil semantic concepts represented by unstructured data instances. Vermeer *et al.* (2019) utilized machine learning to analyze consumer opinions and experiences shared on social media platforms like Facebook and Twitter. Timoshenko and Hauser (2019) employed NLP to extract informative and non-repetitive sentences for subsequent human assessment of customer needs. Their automated UGC analysis method reduced costs by 46-52% compared to traditional qualitative data collection and analysis. Dzyabura and Peres (2021) applied machine learning to capture consumer perceptions of brands by analyzing semantic tags assigned to photos using automated topic detection methods. Hartmann *et al.* (2021) developed a deep-learning pipeline using a neural network pre-trained on 1.2 million images from ImageNet to automatically classify image and text content into categories like "pack shots", "consumer selfies", and "brand selfies."

4 Proposed Methodology

Artificial intelligence (AI) can serve as an instrument for marketing research, strategy development and action planning (Huang & Rust, 2021). Our methodology leverages AI to help customer experience management (CEM) understand topics in UGC. By doing so, CEM can address these topics on the touchpoints under the company's control, thereby creating a consistent brand image.

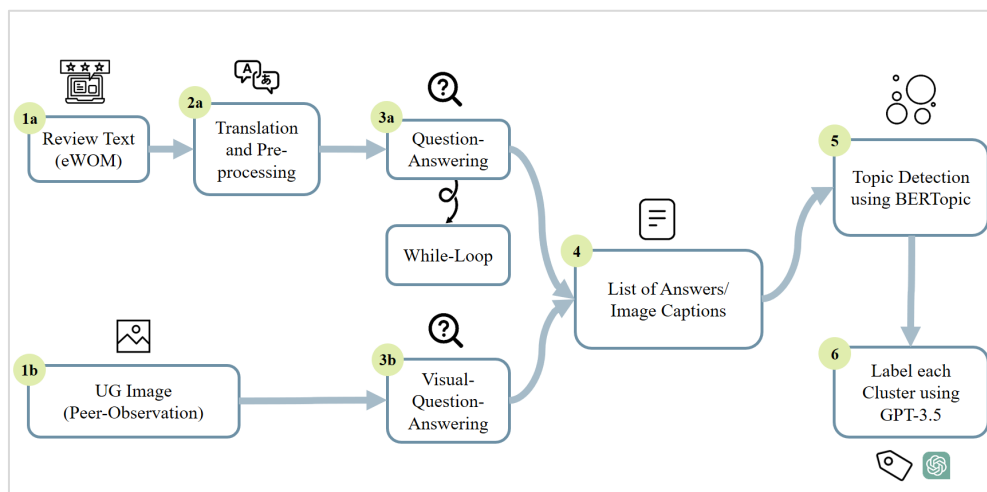


Figure 1: Analysis Pipeline for Automated Topic Extraction from Text and Image Data

Figure 1 illustrates our AI-enabled framework for automatic topic identification, focusing on eWOM in form of Amazon reviews as well as on peer-observable content in form of Instagram posts. To analyze text content, NLP tools are directly applied on the reviews. For image-based UGC analysis, a text caption is first generated using AI-powered methods and then fed into the NLP-analysis pipeline.

4.1 Analyzing eWOM with NLP

NLP involves computer-based techniques for comprehending human language, as outlined by Hirschberg and Manning (2015). A relatively new area within NLP is its generative component, which has been significantly advanced by emerging large-language models (LLMs) like GPT-3.5. These models provide an interface for real-time conversational human-AI interaction, fostering developments in AI (Jurafsky & Martin, 2023; OpenAI, 2023). NLP applications are diverse, ranging from detecting specific words or names within a text to analyzing sentiments and their context. A pivotal milestone in NLP research is the bi-directional encoder representations from transformers (BERT) model by Devlin et al. (2019). BERT is pre-trained on a vast set of non-annotated texts, encompassing over 3 billion words, enabling it to accurately capture semantic meaning within word contexts. In our methodology, we employ topic clustering approaches based on BERT, such as BERTopic, which are particularly effective for analyzing text-based UGC. These approaches automatically handle data sources with large amounts of irrelevant information, such as typos or stop-words (Grootendorst, 2022), making them ideal for extracting meaningful insights from UGC.

After generating a sufficient database of relevant reviews (1a), several steps of pre-processing are applied to normalize and remove the noisy text data instances. The next step involves the translation of texts into English using the NLLB model (Costa-jussà et al., 2022) to unify the data and facilitate further analysis (2a). The NLLB model offers a pre-trained approach for neural machine translation (NMT), built using a transformer-based architecture. Subsequently, the methodology employs a question-answering-procedure (QA) (3a), where GPT-3.5 is used to iteratively apply questions (appendix A) on each review, targeted at extracting only product-related information. As one review can potentially contain a multitude of different brand- or product-related information, the QA-based aspect detection procedure is applied in a loop, until no new product-related statements can be found. As a result, a list of product-related aspects is retrieved from the product reviews, which can be used for the sub-sequent clustering procedure, based on BERTopic (4). We convert these answers into high-dimensional vectors, with one vector representing an answer, using the python sentence-transformer library (Reimers & Gurevych, 2019). The dimensionality of these vectors is optimized using the Uniform Manifold Approximation and Projection (UMAP) algorithm (McInnes et al., 2020). We then apply HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise) (Malzer & Baum, 2020) to the reduced vectors to aggregate the individual text-vectors into cohesive groups (5). This clustering method groups semantically similar instances and organizes them into detailed topic clusters. Finally, we use the GPT-3.5 model to generate representative titles for each cluster, summarizing the core content of the texts within each group based on text samples handed over to the GPT-model (6). These cluster titles can be understood as product-related topics, repeatedly discussed by customers, thus indicating a high relevance for the CEM.

4.2 Analyzing Peer-Observation with IR and NLP

To analyze UGC in the form of Instagram posts we employ IR techniques. IR can generally be understood as machine-learning based methodology for extracting and analyzing visual content from unstructured image data, capturing the contained semantic information and embedding it in the respective context. In the presented work, visual question answering (VQA) plays an important role for the extraction of brand- and product related aspects. OFASys is a versatile model learning system used for various machine learning tasks, including image generation and video captioning (Bai *et al.*, 2022). The initial training of OFASys involved the VQA dataset proposed by Agrawal *et al.* (2015), which comprises images depicting common objects in context (MS COCO dataset) (Lin *et al.*, 2014), along with a dataset of manually labeled descriptive captions for each image. These image-caption-pairs were used to train and validate the question-answering algorithm that can be applied to unlabeled images. The approach enables the analysis of visual UGC to identify relevant objects and colors.

After identifying and collecting relevant posts from social media (1b), the images are entered into an analysis pipeline, designed in a similar way as the review mining. In the context of the proposed framework, VQA can be applied to extract product-related features, use-cases and other aspects of relevance for CE analysis and CEM (3b). The questions utilized for extracting information from visuals can be found in appendix A. As a result, each image is assigned one descriptive caption. After representing visuals through text answers (4), a conventional BERTopic approach is employed on the obtained image captions. Hence, text embeddings are created using sentence-transformers, reduced in their dimensionality through UMAP and clustered with HDBSCAN (5), followed by title generation through GPT-3.5 (6).

5 Use-Case for Methodology Application

The proposed methodology is tested using two main data sources: Amazon reviews, representing eWOM, as well as Instagram images, representing instances of peer-observation (examples can be found in appendix B). The Amazon dataset includes $N=691$ FIMOair reviews from 06/08/2010 to 06/01/2022. After pre-processing, the texts undergo a QA-process, generating $N_{\text{Answers}}=2,313$ answers (product-related aspects) with an average of three answers per review. The minimum cluster size is adjusted based on the number of relevant instances, resulting in $C=30$ clusters. On average, each cluster contains $N_{\text{average_cluster_size}}=77$ answers, with a maximum of $N_{\text{maximum_cluster_size}}=176$ instances. The most prominent topics in the dataset are "*Texture and Material*" ($N=176$) and "*Hand and Foot Prints*" ($N=172$).

To create a dataset of FIMOair related photos, all posts accompanied by the hashtag #fimoair or the combination of hashtags #fimo and #airdryclay are queried and extracted from Instagram. This approach results in $N=3,056$ posts uploaded between 03/15/2013 and 01/16/2023. Here, the clustering procedure yields $C=20$ clusters, with an average number of $N_{\text{average_cluster_size}}=152$ answers inside each cluster and a maximum of $N_{\text{maximum_cluster_size}}=593$ instances. "*Toy Figurines and Stuffed Animals*" ($N=593$), followed by "*Jewelry and Ornaments*" ($N=561$) are the largest clusters. All retrieved topics can be found in appendices C and D.

5.1 Methodology Evaluation

A random sample of 100 Amazon reviews was manually analyzed to extract product-related aspects, and the results were compared to automatically generated QA results, demonstrating the algorithm's ability to accurately detect 66.4% of human-identified information. The algorithm was also capable to detect aspects which were missed by the human annotator, resulting in an overall accuracy of 69.6%. The performance of the VQA algorithm was evaluated by comparing auto-generated image captions with the respective image and deciding whether the caption accurately represents its content. The evaluation was performed using 100 randomly selected auto-annotated pictures, resulting in a 72.0% accuracy.

To assess the quality of the clustering algorithm, three clusters (small, medium and large) were selected from both, the review and image datasets. The individual texts (QA- and VQA-output) within each cluster were compared to the assigned label, determining the cluster accuracy. For instance, the text *"really good quality and great price"* is represented correctly through the cluster *"price and quality"*, while the statement *"we used this product in school"* is semantically misallocated. As a result, the image caption clustering achieves 77.9% to 84.0% accuracy, with the best performance in the smallest cluster ($N_{\text{large}}=243$, $p_{\text{correct}}=81.8\%$; $N_{\text{medium}}=104$, $p_{\text{correct}}=77.9\%$; $N_{\text{small}}=50$, $p_{\text{correct}}=84.0\%$). For review-text-based clustering, the highest accuracy is observed in the medium-sized cluster ($N_{\text{large}}=110$, $p_{\text{correct}}=81.2\%$; $N_{\text{medium}}=76$, $p_{\text{correct}}=88.1\%$; $N_{\text{small}}=24$, $p_{\text{correct}}=62.5\%$) with up to 88.1% accuracy.

Discussion and Further Research

This study aims to develop and validate a methodological framework for marketing service practitioners to create a consistent consumer experience by summarizing topics from earned consumer-brand touchpoints. Consequently, marketers should incorporate the identified topics into their own marketing channels, facilitating a cohesive CE, by delivering inspiration for use cases (e.g., *"Jewelry and Ornaments"*) or by offering solutions to consumer's questions and challenges (e.g., *"Cracks and Brittleness"*). The framework utilizes NLP and IR techniques to extract brand-related insights from unstructured text and images, specifically UGC. In conclusion, the incorporation of NLP and IR into the CEM workflows can offer significant support to market researchers and corporate actors in strategic and operational decision-making, extracting consumer-driven topics for understanding customers and shaping brand-owned communication channels. The findings of this study demonstrate the effectiveness of transformer-based methods for automatic aspect detection and image captioning, yielding satisfactory results at approximately 70 percent accuracy. The results were obtained using BERT-based models and GPT-3.5, suggesting the need for further exploration and experimentation with advanced language models such as GPT4o, which are capable of analyzing images. Future practical efforts should focus on further the automation of the marketing activities, incorporating consumer-relevant topics from earned channels into owned and paid channel content.

References

- Agrawal, A., Lu, J., Antol, S., Mitchell, M., Zitnick, C.L. & Batra, D. et al. (2015) *VQA: Visual Question Answering*. Available from: <https://doi.org/10.48550/arXiv.1505.00468>.
- Babić Rosario, A., Valck, K. de & Sotgiu, F. (2020) Conceptualizing the electronic word-of-mouth process: What we know and need to know about eWOM creation, exposure, and evaluation. *Journal of the Academy of Marketing Science*, 48(3), 422–448. Available from: <https://doi.org/10.1007/s11747-019-00706-1>.
- Bai, J., Men, R., Yang, H., Ren, X., Dang, K. & Zhang, Y. et al. (2022) *OFASys: A Multi-Modal Multi-Task Learning System for Building Generalist Models*. Available from: <https://doi.org/10.48550/arXiv.2212.04408>.
- Balducci, B. & Marinova, D. (2018) Unstructured data in marketing. *Journal of the Academy of Marketing Science*, 46, 557–590. Available from: <https://doi.org/10.1007/s11747-018-0581-x>.
- Barwitz, N. & Maas, P. (2018) Understanding the Omnichannel Customer Journey: Determinants of Interaction Choice. *Journal of Interactive Marketing*, 43, 116–133. Available from: <https://doi.org/10.1016/j.intmar.2018.02.001>.
- Baxendale, S., Macdonald, E.K. & Wilson, H.N. (2015) The Impact of Different Touchpoints on Brand Consideration. *Journal of Retailing*, 91(2), 235–253. Available from: <https://doi.org/10.1016/j.jretai.2014.12.008>.
- Berger, J., Humphreys, A., Ludwig, S., Moe, W.W., Netzer, O. & Schweidel, D.A. (2020) Uniting the Tribes: Using Text for Marketing Insight. *Journal of Marketing*, 84(1), 1–25. Available from: <https://doi.org/10.1177/0022242919873106>.
- Colicev, A., Kumar, A. & O'Connor, P. (2019) Modeling the relationship between firm and user generated content and the stages of the marketing funnel. *International Journal of Research in Marketing*, 36, 100–116. Available from: <https://doi.org/10.1016/j.ijresmar.2018.09.005>.
- Costa-jussà, M.R., Cross, J., Çelebi, O., Elbayad, M., Heafield, K. & Heffernan, K. et al. (2022) *No Language Left Behind: Scaling Human-Centered Machine Translation*. Available from: <https://doi.org/10.48550/arXiv.2207.04672>.
- Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. (2019) BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics: Minneapolis, Minnesota, pp. 4171–4186.
- Dzyabura, D. & Peres, R. (2021) Visual Elicitation of Brand Perception. *Journal of Marketing*, 85(4), 44–66. Available from: <https://doi.org/10.1177/0022242921996661>.
- Grootendorst, M. (2022) *BERTopic: Neural topic modeling with a class-based TF-IDF procedure*. Available from: <https://doi.org/10.48550/arXiv.2203.05794>.
- Hartmann, J., Heitmann, M., Schamp, C. & Netzer, O. (2021) The Power of Brand Selfies. *Journal of Marketing Research*, 58(6), 1159–1177. Available from: <https://doi.org/10.1177/00222437211037258>.
- Hirschberg, J. & Manning, C.D. (2015) Advances in natural language processing. *Science*, 349(6245), 261–266. Available from: <https://doi.org/10.1126/science.aaa8685>.
- Homburg, C., Jozić, D. & Kuehnl, C. (2017) Customer experience management: toward implementing an evolving marketing concept. *Journal of the Academy of Marketing Science*, 45(3), 377–401. Available from: <https://doi.org/10.1007/s11747-015-0460-7>.

- Hoyer, W.D., Kroschke, M., Schmitt, B., Kraume, K. & Shankar, V. (2020) Transforming the Customer Experience Through New Technologies. *Journal of Interactive Marketing*, 51, 57–71. Available from: <https://doi.org/10.1016/j.intmar.2020.04.001>.
- Huang, M.-H. & Rust, R.T. (2021) A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, 49(1), 30–50. Available from: <https://doi.org/10.1007/s11747-020-00749-9>.
- Jurafsky, D. & Martin, J.H. (2023) *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, 3rd edition.
- Keller, K.L. (2013) *Strategic Brand Management: Building, Measuring, and Managing Brand Equity*, 4th edition. Pearson Education Limited: Harlow.
- Keyser, A. de, Verleye, K., Lemon, K.N., Keiningham, T.L. & Klaus, P. (2020) Moving the Customer Experience Field Forward: Introducing the Touchpoints, Context, Qualities (TCQ) Nomenclature. *Journal of Service Research*, 23(4), 433–455. Available from: <https://doi.org/10.1177/1094670520928390>.
- Kuehnl, C., Jozic, D. & Homburg, C. (2019) Effective customer journey design: consumers' conception, measurement, and consequences. *Journal of the Academy of Marketing Science*, 47(3), 551–568. Available from: <https://doi.org/10.1007/s11747-018-00625-7>.
- Lemon, K.N. & Verhoef, P.C. (2016) Understanding Customer Experience Throughout the Customer Journey. *Journal of Marketing*, 80(6), 69–96. Available from: <https://doi.org/10.1509/jm.15.0420>.
- Lin, T.-Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R. & Hays, J. et al. (2014) *Microsoft COCO: Common Objects in Context*. Available from: <https://doi.org/10.48550/arXiv.1405.0312>.
- Liu, L., Dzyabura, D. & Mizik, N. (2020) Visual Listening In: Extracting Brand Image Portrayed on Social Media. *Marketing Science*, 39(4), 669–686. Available from: <https://doi.org/10.1287/mksc.2020.1226>.
- Malzer, C. & Baum, M. (2020) A Hybrid Approach To Hierarchical Density-based Cluster Selection. In: *2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI), 2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, 9/14/2020 - 9/16/2020, Karlsruhe, Germany. IEEE: [S.l.], pp. 223–228.
- McInnes, L., Healy, J. & Melville, J. (2020) *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction*. Available from: <https://doi.org/10.48550/arXiv.1802.03426>.
- Nam, H. & Kannan, P.K. (2020) Digital Environment in Global Markets: Cross-Cultural Implications for Evolving Customer Journeys. *Journal of International Marketing*, 28(1), 28–47. Available from: <https://doi.org/10.1177/1069031X19898767>.
- OpenAI (2023) *Introducing ChatGPT*. Available from: <https://openai.com/blog/chatgpt> [Accessed 28 February 2023].
- Reimers, N. & Gurevych, I. (2019) *Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks*. Available from: <https://doi.org/10.48550/arXiv.1908.10084>.
- Stephen, A.T. & Galak, J. (2012) The Effects of Traditional and Social Earned Media on Sales: A Study of a Microlending Marketplace. *Journal of Marketing Research*, 49(5), 624–639. Available from: <https://doi.org/10.1509/jmr.09.0401>.

- Timoshenko, A. & Hauser, J.R. (2019) Identifying Customer Needs from User-Generated Content. *Marketing Science*, 38(1), 1–20. Available from: <https://doi.org/10.1287/mksc.2018.1123>.
- Tirunillai, S. & Tellis, G.J. (2012) Does Chatter Really Matter? Dynamics of User-Generated Content and Stock Performance. *Marketing Science*, 31(2), 198–215. Available from: <https://doi.org/10.1287/mksc.1110.0682>.
- Vermeer, S.A., Araujo, T., Bernritter, S.F. & van Noort, G. (2019) Seeing the wood for the trees: How machine learning can help firms in identifying relevant electronic word-of-mouth in social media. *International Journal of Research in Marketing*, 36(3), 492–508. Available from: <https://doi.org/10.1016/j.ijresmar.2019.01.010>.
- Wang, Y. & Yu, C. (2017) Social interaction-based consumer decision-making model in social commerce: The role of word of mouth and observational learning. *International Journal of Information Management*, 37(3) 179-189. Available from: <https://doi.org/10.1016/j.ijinfomgt.2015.11.005>.
- Wolny, J. & Charoensuksai, N. (2014) Mapping customer journeys in multichannel decision-making. *Journal of Direct, Data and Digital Marketing Practice*, 15(4), 317–326. Available from: <https://doi.org/10.1057/dddmp.2014.24>.

Appendix A: Questions Defined for Text-QA and Image-QA

The following questions were defined for the goal of deriving product-related topics from unstructured data, such as text and images. In an iterative process, all defined questions are applied to the text entities, resulting in a list of answers. For the VQA task, the questions address the context of the image in terms of objects or concepts.

Text-Data

- *What are clay characteristics?*
- *What are clay features?*
- *What are clay aspects?*

- *What are product characteristics?*
- *What are product features?*
- *What are product aspects?*

Image-Data

- *What does the image describe?*

Appendix B: Exemplary FIMOair Reviews and Instagram Images

Amazon Reviews (translated to English)

1. *"I bought this item for my sister's birthday and she was very happy with it. She makes her own figures and 3D pictures with it. Especially the lightness of the material, she says, is brilliant!"*
2. *"It dries quickly but it crumbles and sticks a little everywhere."*
3. *"Really good quality and great price. Easy to mold and use, my daughter has found this easy to use and loves the color of this clay. Comes packaged really well."*
4. *"Super model mass, there are some in discount stores for half the price, but they can't match the quality, smooth results, easy to process."*
5. *"We bought it for the first time to make Christmas tree ornaments that were light and not easily broken. It did its job very well and we repeated it for other projects with children (5-9 years old). It doesn't stain much and is easy to clean. It hardens in the microwave. The only drawback is that being plastic, it can only be painted with acrylics, and they have to be of good quality, otherwise they don't catch the pigment well. I really liked the discovery."*

Instagram Images




Figure 2: Exemplary Instagram images with FIMOair-related hashtags

Appendix C: List of Topics from Amazon Reviews

Topic Label	Number of Aspects
Texture and Material	176
Hand and foot prints	172
Color (specifically focusing on various shades of white)	153
Ease of Work/Handling	137
Fimo Air Light	134
Breakable items	112
Cracks and brittleness in products	111
Delivery and Service	110
Softness and Malleability	97
Odor/Smell	85
Goodness/Quality	83
Working with kids and projects using Fimo Air Natural	82
Surface Smoothness	77
Shaping and Molding	71
Food processing efficiency	69
Product Durability	66
Product satisfaction	63
Modeling	48
Dough Processing	48
Paste for Modeling/Molding	47
Cooking/ Baking	44
Holiday and Festive Decorations	44
Drying Time	41
Price and Quality	41
Arts and Crafts Supplies	41
Product versatility/Usability	35
Ease of Use	35
Painting and After-treatment	33
Polymer clay products	30
Drying Speed	28

Appendix D: List of Topics from Instagram posts

Topic Label	Number of Posts
Toy figurines and stuffed animals	593
Jewelry and ornaments	561
Flower arrangements and cups	328
Office and art supplies	183
Cake Designs and Toppers	141
Cutting and Crafting Supplies	128
Cookies and ice cream treats	120
Dessert Plates and Food Faces	117
Plates, Bowls, and Wooden Dishware	105
Home Decor with Candles and Coasters	104
Keychains and Charms	97
Statues and Figurines	96
Cupcake Varieties and Decorations	87
Skull and Creature Sculptures	82
Bird-related Decor and Figurines	81
Cactus and Plant Pots	69
Shells and Rocks	51
Easter Eggs and Baskets	51
Mushroom Houses and Gardens	49
Strange and Nonsensical Objects	13