

# Tell Me What Is Good About This Property: Leveraging Reviews For Segment-Personalized Image Collection Summarization

Monika Wysoczanska<sup>1</sup>, Moran Beladev<sup>2</sup>, Karen Lastmann Assaraf<sup>2</sup>, Fengjun Wang<sup>2</sup>, Ofri Kleinfeld<sup>2</sup>, Gil Amsalem<sup>2</sup>, Hadas Harush Boker<sup>2</sup>

<sup>1</sup>Warsaw University of Technology, Warsaw, Poland

<sup>2</sup>Booking.com, Tel Aviv, Israel

monika.wysoczanska.dokt@pw.edu.pl

## Abstract

Image collection summarization techniques aim to present a compact representation of an image gallery through a carefully selected subset of images that captures its semantic content. When it comes to web content, however, the ideal selection can vary based on the user’s specific intentions and preferences. This is particularly relevant at Booking.com, where presenting properties and their visual summaries that align with users’ expectations is crucial. To address this challenge, we consider user intentions in the summarization of property visuals by analyzing property reviews and extracting the most significant aspects mentioned by users. By incorporating the insights from reviews in our visual summaries, we enhance the summaries by presenting the relevant content to a user. Moreover, we achieve it without the need for costly annotations. Our experiments, including human perceptual studies, demonstrate the superiority of our cross-modal approach, which we coin as *CrossSummarizer* over the no-personalization and image-based clustering baselines.

## Introduction

Visual content is one of the key aspects when evaluating and deciding upon a place to stay on the Booking.com platform. Throughout their journey, platform users browse through visual content for four main reasons: (1) To get an accurate and realistic idea of what to expect, (2) To assess the quality of the property, (3) To build trust and remove doubts that they are making the right booking decision, (4) To look for a content that matches their travel intent. When looking for their next trip, users might be overwhelmed with the amount of information they are exposed to, both visual and textual. Image galleries can contain up to hundreds of images. Hence, we aim to focus the platform’s users on the visual content which is the most relevant to them given their current personal context. We achieve that by summarizing each property with a subset of visually informative, high-quality, segment-personalized images.

Most of the works in the image collection summarization area focus on a generic summarization problem, where the main objective is to select a *diverse* set of images. Only recently, some of the efforts have been made in a so-called *guided summarization* (Kothawade et al. 2022), where the

aim is to get a diverse yet *representative* subset of images corresponding to a specific query. In our work, personalization can be seen as a variant of a query-based approach. However, the queries are not explicit. User intents can not easily be translated into specific queries, making *personalized image gallery summarization* more complex. In this work, inspired by recent advances in multi-modal learning, we solve the above-mentioned challenge and develop a method for personalized image collection summarization with textual guidance. We focus on entire groups of users, which we further refer to as **user segments**. We leverage millions of reviews corresponding to properties on the Booking.com platform for enhanced user segment personalization. We do that by extracting key topics mentioned in the reviews, and since they significantly differ across segments of users (see Fig. 1) we adjust the image selection accordingly. We identify the following challenges for our task:

**Personalization modelling** It is not apparent how to obtain the personalization data for our task and avoid costly annotations. Therefore, we focus on the entire segments of users and leverage the textual reviews with the metadata available on the Booking.com platform.

**User intents extraction from reviews** Reviews available on the platform can be a very rich source of information. Users of the platform can help future travellers to make the best choices by sharing their experiences. However, in practice, the reviews tend to be very noisy. Therefore, an essential aspect of our approach is to extract relevant pieces of information for a visual summary from free-form text. Thus, the extracted signal has to be adequate for matching the semantic content of images in a gallery.

**Matching text and images** Finally, matching text with images requires representing all of them in a joint multi-modal space. There has been a recent surge of methods that leverage free-form text to reduce the need for costly annotations (Narasimhan et al. 2022; Wang et al. 2023; Stroud et al. 2020) yielding better performances when cast into multi-modal problems (Li et al. 2022; Narasimhan, Rohrbach, and Darrell 2021; Zhang et al. 2022; Zhu et al. 2018, 2020; Wysoczanska. and Trzciński. 2020), or obtaining better image representations (Morgado, Vasconcelos, and Misra 2021; Radford et al. 2021). Most importantly, CLIP (Radford et al. 2021) shows that image-text large-scale

pretraining gives the ability to learn the generalizable image representations and enables zero-shot image-text matching, which we leverage in our approach.




Figure 1: Heatmap of most popular topics (x-axis) extracted from reviews for different traveller types at Booking.com. We note that the ranking of the most mentioned topics differs among traveller segments making reviews a valuable signal for user segment personalization.

The main contribution of our work is as follows:

- We introduce an unsupervised method for image collection summarization personalized for entire segments of users using textual guidance extracted from reviews. Our approach leverages text and image representations in multi-modal space,
- We extend previously introduced evaluation metrics for image collection summarization to the segment-personalization use case when no ground-truth annotations are given,
- In our experiments, we conduct human perceptual studies alongside the quantitative evaluation using our proposed metrics and show that reviews provide an adequate signal for segment personalization.

## Method

In this section, we define a task of personalized image collection summarization for segments of users and describe in detail all the building blocks of *CrossSummarizer*. We also provide the metrics which we use to evaluate our method quantitatively. We explain how we adapt standard metrics for image collection summarization to our segment-personalized use case.

### Task Definition

Given a collection of images  $G$ , the goal of our method is to select a subset of  $K$  images  $G_s \subset G$  that best corresponds to a given user segment  $u \in U$ . We consider two types of user segments:

- Travellers: *Solo, Couple, Group, Family, Business*
- Trips: *Beach, Ski, City, Nature Active, Nature Peaceful*

Our personalized summary aims to cover the essential aspects of the whole image gallery while selecting images that show details relevant to the specified user segment  $u$ .

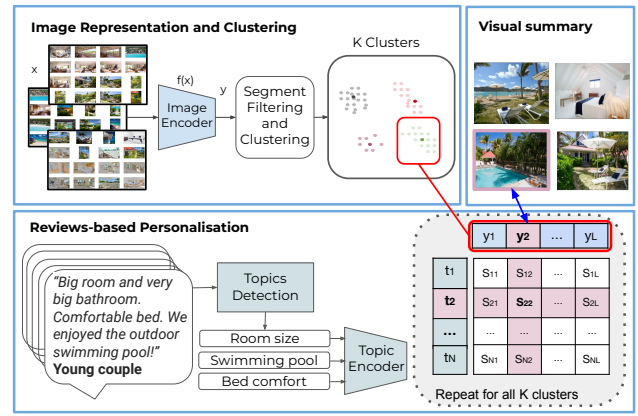


Figure 2: Overview of our method. First, we extract image embeddings and cluster them to obtain  $K$  ( $K=4$  in this case) semantically separated groups of images. Then, for each cluster, we calculate the similarities between all the images within the cluster and the topics extracted from reviews of the specific segment  $u$  (here  $u=Couple$ ). Finally, the selected images are the ones with the highest similarity to any of the topics.

Each of the images in  $G$  has a corresponding set of visual classes  $c \in C$ , and for each user segment  $u$ , we manually define a subset of  $C_u \subset C$  of relevant classes. Moreover, we have access to each property’s textual reviews, labelled with the corresponding segment  $u$ .

However, reviews can mention many aspects in just a short piece of text (see Fig. 2). Therefore, we represent each review as a set of topics  $T$  it covers.

Below we explain all of the steps of our method in detail.

### Image Embeddings & Filtering

Let  $f(x_i) \rightarrow y_i \in \mathbb{R}^D$  be a feature representation of an image  $x_i \in G$ . In this work, we use a recently proposed MuMIC (Wang et al. 2023) for image representation  $f(x)$ , which extends the image-text contrastive pretraining proposed in CLIP (Radford et al. 2021) to a multi-label case. MuMIC learns a multi-modal embedding space by jointly training an image encoder and text encoder to maximize the cosine similarity of the image and label-text embeddings of real labels. The model applies tempered sigmoid-based Binary Cross-Entropy (BCE) loss on each class and then mean-reduces it (see Eq. 1), optimizing across all classes. Given a batch of images  $\{x_i \in G, i = 1..N\}$  and their associated ground-truth multi-label vector  $\{\hat{w}_i \in \mathbb{R}^{|C|}, i = 1..N\}$

$$\ell_{BCE} = -\frac{1}{N|C|} \sum_{i=1}^N \sum_{j=1}^{|C|} (p_j \hat{w}_{ij}) \cdot \log \sigma(w_{ij}) + (1 - \hat{w}_{ij}) \cdot \log(1 - \sigma(w_{ij})) \quad (1)$$

where  $\sigma(\cdot)$  is the Tempered Sigmoid function;  $w_{ij}$  are the original output logits (before applying temperature scaling), which is the pairwise image-text cosine similarity between image  $x_i$  and class  $j$ ; and  $p_j$  is the positive sample weight of

class  $j$ . A higher  $p_j$  indicates that positive samples are given greater weight, increasing the penalty for identifying false negatives.

We use MuMIC to both extract image embedding as well as to obtain classes  $C$  per each image  $x \in G$ . We then filter out images that are not relevant to given  $u$  by keeping only the subset of images with at least one class  $c \in C_u$ .

## Clustering

For the clustering phase, we use *KMedoids* algorithm (Park and Jun 2009) implementation. We choose *KMedoids* as it is robust to outliers and gives high flexibility in choosing  $K$ .

We run clustering on the image embeddings using *Cosine Similarity* (*CosSim*) as a similarity metric, such that

$$\text{CosSim}(u, v) = \langle \frac{u}{\|u\|}, \frac{v}{\|v\|} \rangle, \quad \text{with } u, v \in \mathbb{R}^D \quad (2)$$

where  $\langle \cdot, \cdot \rangle$  is the inner product in  $\mathbb{R}^D$ .

We thus obtain cluster assignments for each of the images.

## Text2topic: Topics Detection Model

We use a recently proposed Text2Topic (Fengjun Wang, Moran Beladev et al. 2023) model, a topic detection model, to extract user segment preferences and personalize the subset of images based on their topics of interest. We first filter the reviews by  $u$  and then detect the topics associated with these reviews. To get the topics, we train a classification model with 45 travel-domain topics (see the subset of the topics in the heatmap in Fig. 1) using cross-encoder transformer-based architecture (Reimers and Gurevych 2019), which relies on BERT (Devlin et al. 2019). We train the model with 15,663 positive pairs of reviews and topics and sample X5 negative topics per review. We use the Binary Cross Entropy loss function on the [CLS] embedding vector of the crossed input [REVIEW] [SEP] [TOPIC] to get the probability the topic is mentioned in the review.

At inference, we run the model on pairs of reviews and each of the 45 topics to get the probability scores. Finally, we get the topics  $T$  with a probability greater than 0.5 to match with the review.

## Matching Images to Topics

Having obtained clusters of images and a list of topics  $T$  for given  $u$ ,  $T_u \subset T$  we select the final subset of images  $G_s$ . We do it by first computing the *confidence matrix*  $S \in \mathbb{R}^{2 \times 2}$  of images  $G$  being aligned with topics  $T_u$ . We follow (Wang et al. 2023) and use tempered sigmoid, formulated as:

$$S_{ij} = \sigma(\exp(\gamma) \cdot \langle t_i, y_j \rangle) \quad (3)$$

where  $\gamma$  is the log-parameterized multiplicative scalar, and  $t_i, y_j \in \mathbb{R}^D$  are respectively a topic from  $T_u$  and a feature representation  $f(x_j)$  of an image  $x_j$ .

We then select the final images by iterating over clusters and selecting the pair  $(t_i, y_j)$  with the highest similarity within a cluster of representations. The pseudo-code for this selection is given in the Algorithm 1.

---

### Algorithm 1: Pseudocode for selecting images

---

**Result:**  $G_s$  - selected images from the gallery

**Input:**

$A_K$  - Image-to-cluster assignment :  $x \in G \mapsto \{1..K\}$

$S$  - Confidence matrix from Eq. (3)

**procedure** SelectImages ( $L_k, S$ )

```

 $G_s \leftarrow \emptyset$  // Selected images
 $\Omega_T \leftarrow \{1..|T|\}$  // Indices of active topics
for Cluster  $k \in \{1..K\}$  do
    // Image indices for cluster  $k$ 
     $\Omega_k \leftarrow \{j : A_K(x_j) = k, \forall j \in \{1..|G|\}\}$ 

    // Compute best matches within  $k$ 
     $i^*, j^* \leftarrow \underset{(i,j) \in \Omega_T \times \Omega_k}{\operatorname{argmax}} S_{ij}$ 

     $\Omega_T \leftarrow \Omega_T \setminus \{i^*\}$  // Update active topics
     $G_s \leftarrow G_s \cup \{x_{j^*}\}$  // Add selected image
return  $G_s$ 

```

---

## Evaluation Metrics

To ensure our generated summaries are diverse, and adequately correspond to user segments' interests, we define multiple evaluation metrics. Following (Iyer et al. 2018), we use *Coverage*, *Representativeness* and *Diversity* metrics. However, we apply some modifications to match our use case. More specifically, we make sure each one of the metrics is normalized across samples.

**Diversity** Let us denote  $d(x_i, x_j)$  as a distance between images  $x_i$  and  $x_j$ . The *Diversity* (*Div*) metric measures to what extent the diversity in terms of distance between embeddings of images in  $G_s$  is similar to the one in the original gallery  $G$ . We define *Div* as:

$$\text{Div} = \frac{\max_{(x_i, x_j) \in G_s \times G_s} d(x_i, x_j)}{\max_{(x_l, x_m) \in G \times G} d(x_l, x_m)}, \quad (4)$$

where  $d(\cdot, \cdot)$  is computed in image embedding space as:

$$d(x_i, x_j) = 1 - \text{CosSim}(f(x_i), f(x_j)). \quad (5)$$

**Representativeness** Then, let us denote  $\mu_G, \mu_{G_s} \in \mathbb{R}^D$  as the mean vectors of the original gallery representations and the selected subset, respectively. *Representativeness* (*Repr*) is defined as:

$$\text{Repr} = \text{CosSim}(\mu_G, \mu_{G_s}) \quad (6)$$

With the generated summaries, we wish both vectors to be similar in representation, such that their CosineSimilarity is close to 1.

**Coverage** We also measure *Coverage* (*Cov*) in the semantic space by using classes associated with images in  $G$ . We adopt the *Probabilistic coverage* suggested in (Iyer et al. 2018) and use the probabilities  $P(c|x) \in \mathbb{R}$  for a given image  $x$  to represent a particular class  $c \in C_u$ .

Split	No. reviews	No. images	No. samples	Total No. reviews	Total No. images
Small	between 30 and 150	between 50 and 100	3230	251943	232560
Big	151 <	100 <	3151	1504024	457946

Table 1: Dataset split in detail. Overall we collected more than 6000 samples from the platform. Through stratified sampling, we make sure the distribution in terms of location, rating, and accommodation type reflects the real distribution.

$$Cov = \frac{1}{|C_u|} \sum_{c \in C_u} \frac{P_{G_s}(c)}{P_G(c)}, \quad (7)$$

where  $P_G(c) = \max_{x \in G} P(c|x)$ . Note that in our work we use MuMIC model to obtain probabilities since it was trained in-domain. However, this could be any off-the-shelf multi-label image classification method.

With our *Coverage* metric, we mainly focus on measuring the accuracy of the personalization step. Therefore, instead of taking all of the classes in  $C$ , we only consider the ones corresponding to a given user segment  $u$ ,  $C_u$ .

**Reviews Coverage** Finally, to measure how well the generated summaries correspond to user segment topics from reviews, we calculate the topic coverage of selected images. Similarly to *Coverage*, we use the confidence matrix  $S \in \mathbb{R}^2$  as probabilities of a topic  $t_j$  being aligned with the set of selected images  $G_s$ . *Reviews Coverage (RCov)* is therefore given by:

$$RCov = \frac{1}{|T_u|} \sum_{i=1}^{|T_u|} \frac{\max_{j \in \Omega_{G_s}} S_{ij}}{\max_{j \in \Omega} S_{ij}}, \quad (8)$$

where  $\Omega = \{1..|G|\}$  and  $\Omega_{G_s} = \{j : x_j \in G_s, \forall j \in \Omega\}$  are respectively the range of indices of the images in  $G$  and the subset of indices of images that are only in  $G_s \subset G$ .

## Experiments

In this section, we provide the experimental results obtained through both offline evaluation and user studies conducted internally at Booking.com. We also describe our experimental setup, including details on the dataset collected at the Booking.com platform and baseline models we compare our CrossSummarizer against.

### Experimental Setup

**Dataset** We conduct our experiments on the Booking.com dataset consisting of properties. We collected over 6000 real properties from the platform by carefully curating the sampling to adequately represent a real distribution in terms of geographical location, types of accommodation as well as travellers’ experience with a given property. We do it by applying a stratified sampling technique with a country, type of accommodation, and property rating being the factors. We do not share our dataset, however, we note that the examples can easily be downloaded (both images and reviews) since the data is publicly available on the Booking.com site.

Each of the samples in our dataset consists of a set of uploaded by property owners images, which correspond to a

gallery, and a set of reviews of past travellers’ experiences. Alongside sampled reviews, we also include metadata about the type of traveller that authored a particular review.

Since samples in our dataset significantly vary in amounts of both images and reviews, we split the dataset into two groups according to the size of galleries and the number of reviews. Precisely:

- **Small** - properties with the size of a gallery between 50 and 100 photos and 30 - 150 reviews,
- **Big** - gallery sizes of 100 < photos and 150 < reviews.

More details on precise numbers of our dataset split are given in Tab. 1. In our experiment, we report results separately for the two aforementioned splits.

**Baselines** To the best of our knowledge, none of the proposed methods for *image collection summarization* such as (Kothawade et al. 2022; Tschitschek et al. 2014) nor the *multimodal* ones (Zhu et al. 2020; Zhang et al. 2022; Li et al. 2018) tackle the *personalization case*. For evaluation purposes, we implement multiple baselines, which we describe in detail below.

**Default Clustering (Def)** We compare against a simple no-personalization approach by running clustering on image embeddings and omitting the filtering step. The selected  $K$  images are resulting cluster centres. A similar approach was proposed in (Hadi, Essannouni, and Thami 2006) for videos.

**Clustering with personalization (Clust-W/P)** We implement the approach without topic-based refinement and select cluster centres as the summarization. The approach differs from the *Clustering* setting by an additional filtering phase based on relevance to the particular user segment classes being associated with images in the gallery (see Sec. ).

**Topic-based personalization (TopicBased)** We also implement the approach based only on the topic’s similarity with images, without a clustering step. We extract image and topic embeddings and apply user segment filtering. We calculate the similarity matrix between all the image and topic embeddings and choose top  $K$  scores in the matrix. We make sure each image gets selected only once and apply the selection process iteratively. The pseudo-code is similar to the final approach, as we simply take  $\text{argmax}$  on the entire confidence matrix  $S$ , which is being decreased with each iteration by selected images’ columns.

Note that for all of the approaches, we use the same image representations obtained with MuMIC, which was trained on the Booking.com multi-label classification dataset as described in (Wang et al. 2023). For the filtering phase, we create a mapping of MuMIC’s classes for each user segment beforehand and use the mapping at inference time.

## Offline Evaluation

**Evaluation protocol** We first run the evaluation procedure offline internally at Booking.com. For all the methods relying on *K*Medoids algorithm, we set the same random seed for a fair comparison. We run experiments with a fixed  $K = 9$  corresponding to the current Booking.com setting.

**Quantitative results** The results of our experiments are presented in Tab. 2. Looking at *Cov* and *RCov*, we observe that our cross-modal approach outperforms the baselines in the segment personalization task, especially for the **Small** dataset split. Compared to the Def setting, we observe approximately 0.1 gains in *Cov* for both dataset splits and a significant improvement over *RCov*. For the **Big** dataset split, we report a higher *RCov* for the TopicBased approach. This is expected since the approach is based on maximizing the similarity between topics and images, following Eq. 8.

Moreover, our cross-modal approach outperforms the rest of the methods in terms of *Diversity*, indicating that the personalization step produces more diverse summarization.

Relying only on reviews, which corresponds to the results of a Topic-based method, gives a significantly lower *Representativeness* and poor visual user segment *Coverage*. *Coverage* gain of our CrossSummarizer over the Clust-W/P approach also emphasizes the need for using two modalities in our task. Overall, we observe a clear trade-off between *Diversity* and *Representativeness*. CrossSummarizer finds a sweet spot between the two, giving an excellent segment personalization result at the same time.

**Qualitative results** We complement the quantitative evaluation with qualitative results. Fig. 4 shows the visual comparison between the two baseline approaches and our method. The example shows a non-personalized Def result and the personalized results obtained with Clust-W/P and CrossSummarizer approach for *Ski* trip and  $K = 8$ .

All of the approaches give a set of diverse images without any redundant photos. However, our cross-modal model provides the best personalization result. We highlight the photos that are relevant to the user segment: a picture of a skier and a photo of a sauna. We can also see that contrary to the Clust-W/P approach, with our CrossSummarizer, the selected image of a property from outside (bottom right) was taken in the wintertime (expected for *Ski* trip type).

## User Studies

**Experimental setup** In addition to offline evaluation, we conduct human perceptual studies in the form of an anonymized paired test. We compare two methods: Clust-W/P approach and CrossSummarizer. We manually select 210 samples of properties with galleries that contain photos relevant to the specified user segment. We split the samples across 5 participants. The samples are uniformly distributed over the dataset split (**Small**, **Big**) and user segment types.

The participants were asked to answer the question: *Which of the models (A or B) gives a better summary for a given user segment?* We assign a score of 1 for a model that performs better than the other and 0 otherwise. We also allow for a 0.5 score in case of a tie.




Figure 3: Results of our user studies. We report the results separately for each user segment (x-axis).

**Results** The conducted user studies indicate the superiority of our CrossSummarizer approach with an average score of  $0.66 \pm 0.38$  for all user segments over Clust-W/P  $0.34 \pm 0.38$ . Additionally, in Fig. 3, we provide detailed results for each user segment. We observe that CrossSummarizer obtains higher average scores for most of the segments in our experiments, except for only *Business* traveller type. To confirm the statistical significance of our studies, we perform the paired T-test (Student 1908) with a null hypothesis of *average scores for models A and B being equal* and the alternative hypothesis of *model A scoring lower than model B*. The results ( $t(209) = -5.537; p < 5e - 8$ ) let us reject the null hypothesis and accept the alternative hypothesis.

## Application

This section covers some practical aspects of our approach and how it is leveraged at Booking.com. Two essential parts are first run offline, which are the Text2Topic model for topic extraction and the MuMIC model for image embedding extraction and image multi-class annotations. The results are then stored in the database and accessed at runtime. It takes approximately 170 ms for the MuMIC model to run on a batch of 100 images and 244 ms for the Text2Topic model to run prediction on a batch of 100 reviews. We leverage GPU computation for this purpose.

## Deployment & Maintenance

The model is served with Amazon SageMaker and deployed on the Booking.com Content Intelligence Platform (CIP) (Wang et al. 2023). CIP is a stream processing platform based on Apache Flink. It consumes real-time events from Kafka topics (e.g. images uploaded by Booking.com partners) and generates model-based predictions. The same architectural design allows CIP to be used for backfilling purposes. Backfilling refers to the enrichment of historical data with newly deployed model predictions. We leverage the mentioned design by simulating events of historical data and pushing them to Kafka.

CIP is designed to achieve high prediction throughput while keeping a low latency. It achieves that by leveraging

Method	Small				Big			
	Div $\uparrow$	Repr $\uparrow$	Cov $\uparrow$	RCov $\uparrow$	Div $\uparrow$	Repr $\uparrow$	Cov $\uparrow$	RCov $\uparrow$
<i>No Personalization</i>								
Default Clustering (Hadi, Essannouni, and Thami 2006)	0.947	<b>0.931</b>	0.516	0.581	0.925	<b>0.929</b>	0.430	0.488
<i>Personalization</i>								
Clustering with Personalization	0.903	<b>0.931</b>	0.578	0.656	0.932	0.928	0.487	0.521
Topic-Based Personalization	0.903	0.733	0.365	0.649	0.874	0.775	0.261	<b>0.707</b>
CrossSummarizer (Cross-modal)	<b>0.963</b>	0.903	<b>0.617</b>	<b>0.729</b>	<b>0.950</b>	0.885	<b>0.524</b>	0.677

Table 2: Quantitative evaluation. We report the results for two splits: small and big galleries and the number of reviews.

Apache Flink’s asynchronous I/O operator to perform concurrent asynchronous HTTP calls to a model endpoint. However, this optimization mechanism relies on the assumption that each model prediction can be made independently. This assumption does not hold for summarization models where a group of events should be sent to the model together in a single prediction query.

We built upon Apache Flink’s windowing mechanism to implement the grouping of events that should be sent to the model endpoint. Whenever a new call to the model should be triggered, an event containing the request metadata (e.g. a hotel ID) is sent to Kafka. Then, the matching image data is fetched by issuing calls to an external service holding image data. Those calls are executed independently and concurrently. Following that, images are grouped in the same window. Images are accumulated within the window as soon as they are fetched, and the window is closed after a predefined period (e.g. 3 seconds). Then, the images are sent together to the model endpoint for prediction.

### Application Example




Figure 4: Example results of summaries of  $K = 8$  images for one of the properties. We compare Def (top), and personalized methods (bottom) on this property’s visual summarization. The presented personalization examples here are for a *Ski* trip type. We highlight in red the selected images which are relevant to this user segment.

The applications of our image collection summarization approach at Booking.com are three-fold:

- Image subset selection optimization for large collections of images when given a constraint on a number of images or smaller displays,

- Visual content personalization based on traveller type,
- Visual content personalization based on trip type.

Fig. 4 shows some qualitative results of our CrossSummarizer for the third bullet point. Using this personalization approach, we expect to reduce the friction from the decision-making process as users will have a better understanding of the property at an earlier step.

Our proposed model is currently under experimentation. When deploying the personalized CrossSummarizer model we compare it with the current model, which is produced by the Def model, through A/B test experimentation on the CTR (click to ratio) metric.

We also note that personalizing summaries of collections of images is a relevant task for most e-commerce websites. Hence, our method could also be applied to any other personalization task, such as product recommendations where a multi-modal input is available (reviews + images).

### Conclusions & Limitations

We presented a method for personalized image collection summarization for entire segments of users. Our approach is capable of taking into account users’ intents when producing summaries of large image collections. As the personalization signal, we use other travellers’ experiences with properties, which we extract from the reviews. We implemented and tested our method on the Booking.com platform and our experiments, including human perceptual study, indicate that our proposed approach yields good results on the diversity and representativeness axes. The comparison with other baselines indicates that our proposed method performs the best in terms of personalization. Future works include A/B tests of our model in a production environment to measure the real-world impact.

The main limitation of our method is handling samples that are considered a *cold-start* zone, e.g. having a limited number of reviews, but complete image galleries. This, however, can be addressed by leveraging information from other properties of a similar profile.

### Acknowledgments

The work was supported by Booking.com. Monika Wysoczanska was partially supported by the National Centre of Science (Poland) Grant No.2022/45/B/ST6/02817.

## References

- Devlin, J.; Chang, M.-W.; Lee, K.; and 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)*, 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.
- Fengjun Wang, Moran Beladev; Kleinfeld, O.; Frayerman, E.; Shachar, T.; Fainman, E.; Assaraf, K. L.; Mizrachi, S.; and Wang, B. 2023. Text2Topic: Multi-Label Text Classification System for Efficient Topic Detection in User Generated Content with Zero-Shot Capabilities. arXiv:2310.14817.
- Hadi, Y.; Essannouni, F.; and Thami, R. O. H. 2006. Video Summarization by K-Medoid Clustering. In *Proceedings of the 2006 ACM Symposium on Applied Computing, SAC '06*, 1400–1401. New York, NY, USA: Association for Computing Machinery. ISBN 1595931082.
- Iyer, R. K.; Dubal, P.; Dargan, K.; Kothawade, S.; Mahadev, R.; and Kaushal, V. 2018. Vis-DSS: An Open-Source toolkit for Visual Data Selection and Summarization. *CoRR*, abs/1809.08846.
- Kothawade, S.; Kaushal, V.; Ramakrishnan, G.; Bilmes, J.; and Iyer, R. 2022. PRISM: A Rich Class of Parameterized Submodular Information Measures for Guided Data Subset Selection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(9): 10238–10246.
- Li, H.; Ke, Q.; Gong, M.; and Zhang, R. 2022. Video Summarization Based on Video-text Modelling. *CoRR*, abs/2201.02494.
- Li, H.; Zhu, J.; Liu, T.; Zhang, J.; and Zong, C. 2018. Multi-modal Sentence Summarization with Modality Attention and Image Filtering. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, 4152–4158. International Joint Conferences on Artificial Intelligence Organization.
- Morgado, P.; Vasconcelos, N.; and Misra, I. 2021. Audio-visual instance discrimination with cross-modal agreement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12475–12486.
- Narasimhan, M.; Nagrani, A.; Sun, C.; Rubinstein, M.; Darrell, T.; Rohrbach, A.; and Schmid, C. 2022. TL; DW? Summarizing Instructional Videos with Task Relevance and Cross-Modal Saliency. In *European Conference on Computer Vision*, 540–557. Springer.
- Narasimhan, M.; Rohrbach, A.; and Darrell, T. 2021. CLIP-It! Language-Guided Video Summarization. *CoRR*, abs/2107.00650.
- Park, H.-S.; and Jun, C.-H. 2009. A simple and fast algorithm for K-medoids clustering. *Expert Systems with Applications*, 36(2, Part 2): 3336–3341.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Stroud, J. C.; Ross, D. A.; Sun, C.; Deng, J.; Sukthankar, R.; and Schmid, C. 2020. Learning Video Representations from Textual Web Supervision. *CoRR*, abs/2007.14937.
- Student. 1908. The probable error of a mean. *Biometrika*, 1–25.
- Tschiatschek, S.; Iyer, R. K.; Wei, H.; and Bilmes, J. A. 2014. Learning Mixtures of Submodular Functions for Image Collection Summarization. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Wang, F.; Mizrachi, S.; Beladev, M.; Nadav, G.; Amsalem, G.; Assaraf, K. L.; and Boker, H. H. 2023. MuMIC–Multimodal Embedding for Multi-Label Image Classification with Tempered Sigmoid. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 15603–15611.
- Wysoczańska, M.; and Trzcinski, T. 2020. Multimodal Dance Recognition. In *Proceedings of the 15th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 5: VISAPP*, 558–565. INSTICC, SciTePress. ISBN 978-989-758-402-2.
- Zhang, Z.; Meng, X.; Wang, Y.; Jiang, X.; Liu, Q.; and Yang, Z. 2022. UniMS: A Unified Framework for Multimodal Summarization with Knowledge Distillation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10): 11757–11764.
- Zhu, J.; Li, H.; Liu, T.; Zhou, Y.; Zhang, J.; and Zong, C. 2018. MSMO: Multimodal Summarization with Multimodal Output. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4154–4164. Brussels, Belgium: Association for Computational Linguistics.
- Zhu, J.; Zhou, Y.; Zhang, J.; Li, H.; Zong, C.; and Li, C. 2020. Multimodal Summarization with Guidance of Multimodal Reference. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05): 9749–9756.