

# A Baseline Interactive Retrieval Engine for Visual Lifelogs at the NTCIR-14 Lifelog-3 Task

Van-Tu Ninh<sup>1\*</sup>, Tu-Khiem Le<sup>1\*</sup>, Liting Zhou<sup>1</sup>, Graham Healy<sup>1</sup>, Kaushik Venkataraman<sup>1</sup>, Minh-Triet Tran<sup>2</sup>, Duc-Tien Dang-Nguyen<sup>3</sup>, Sinead Smyth<sup>1</sup>, and Cathal Gurrin<sup>1</sup>

<sup>1</sup> Dublin City University, Ireland

<sup>2</sup> VNU-HCM, University of Science, Vietnam

<sup>3</sup> University of Bergen, Norway

**Abstract.** This paper describes the work of DCU research team in collaboration with University of Science, Vietnam, and University of Bergen, Norway at the Lifelog task of NTCIR-14. In this paper, a new interactive retrieval engine is described that supports faceted retrieval and we present the results of an initial experiment with four users. Following this initial experiment, we implement a list of changes for a revised interactive retrieval engine for the LSC2019 comparative evaluation competition. The interactive retrieval system we describe utilises the wide range of lifelog metadata provided by the task organisers to develop an extensive faceted retrieval system.

**Keywords:** Interactive lifelog search engine · Information retrieval

## 1 Introduction

Information Retrieval has a long history of utilising the human as a key component of a retrieval system. Our current generation of WWW search engines rely on the human as an integral part of the search process, in terms of query generation, refinement and result selection. Inspired by the 'human-in-the-loop' model of interactive information retrieval, the DCU team, with the support of VNU-HCM, University of Science and the University of Bergen, developed a prototype interactive retrieval system for the LSAT - Lifelog Semantic Access subtask of the NTCIR-14 Lifelog task [3]. In this paper we introduce this prototype retrieval engine, we present the performance of the retrieval engine in the LSAT task, we report on the findings of a small-scale qualitative user study of the prototype, and we highlight the enhancements carried out on the prototype for our participation in the LSC2019 Lifelog Search Challenge.

## 2 Related Interactive Lifelog Retrieval Systems

The Lifelog Semantic Access Task, which began in NTCIR-12, allows both automatic and interactive lifelog search systems [4] to be comparatively evaluated

---

\* The two authors contributed equally to this paper.

in an open benchmarking exercise. In NTCIR-12, the University of Barcelona and Technical University of Catalonia developed an interactive search engine [16] which utilised visual semantic concepts from images and used them as tags for the interactive image retrieval system. They also employed WordNet to create the similarity between tags to assist novice/expert users to choose the most relevant appropriate tags. Moreover, a heatmap was generated to show the confidence of the retrieval result which aims to achieve the best configuration of precision and recall of their retrieval system. In the official results of the lifelog task, their best run (unsurprisingly) outperforms all the best ones of other teams that built automatic search engines [4]. For the LSAT task at NTCIR-14, [2] developed an interactive lifelog retrieval system that automatically suggested a list of candidate query words to the user and adopted a probabilistic relevance-based ranking function for retrieval. They enhanced the official concept annotations by applying the Google Cloud Vision API<sup>†</sup> and pre-processed the visual content to remove images with poor quality and to offset the fish-eye nature of the wearable camera data. In the provided examples, this was shown to increase the quality of the non-official annotations. Additionally, an interactive system was developed by [15], which operated as a faceted search system over enhanced metadata (additional object detectors, manually adding in ten habit concepts, scene classification, and counting the number of people in the images). The user interface was designed to be user-friendly and support novice users to generate queries and browse results. The system performed significantly better than any other interactive system at NTCIR-14, including the system described in this paper.

More recently, we note the introduction of a new challenge, specifically aimed at comparing approaches to interactive retrieval from lifelog archives. The Lifelog Search Challenge (LSC) [6] utilises a similar dataset [5] to the one used for the NTCIR14-Lifelog task. The LSC has occurred in 2018 and 2019 and attracted significant interest from participants. We report on some of the most relevant of these here. In 2018, six participating teams took part in the live search challenge. These teams had all indexed the dataset prior to attending the workshop and then during the interactive search challenge, both expert and novice users took part in evaluating the performance of the six systems.

Alpen-Adria-Universität Klagenfurt (AAU) developed an interactive retrieval engine based on a video-retrieval system. Called liveXplore [14], a system modification serving as a lifelogging data browser by focusing on visual exploration and retrieval as well as metadata filtering. The system focused on visual similarity, concept and metadata filtering; it performed very well in 2018, coming a very close second place to the eventual winner [1]. A similar version of liveXplore was deployed for the LSC2019 challenge [10], though not as successfully. Another system of note came from Charles University, Prague, with a repurposed an updated version of the VIRET video retrieval system [12]. Every day from the collection was treated as one ‘video’ represented by the lifelog images, with automatic annotations associated with each image using GoogleNet. In ad-

<sup>†</sup> Google Cloud Vision API - <https://cloud.google.com/vision/>

dition, a colour signature for sketch-based search and deep feature vector from the original GoogleNet were extracted. The system came a close third place in both the 2018 and 2019 competitions. Additionally in 2019, we note two additional systems that warrant review. The vitivr system from Rossetto et al. [19] was an enhanced version of the vitivr open-source multimedia retrieval system, which was developed for video retrieval tasks. Extensions to the leading interactive video retrieval system included the capability to process Boolean query expressions alongside content-based query descriptions in order to leverage the structural diversity inherent to lifelog data. This system was the eventual winner of the LSC 2019 competition. A final system of note was developed by [8] which, as per their work at NTCIR-14 [15] included additional enhanced meta-data that proved meaningful for the retrieving process, and a user interface that was designed to support a novice user to perform the retrieval efficiently.

### 3 An Interactive Lifelog Retrieval System

For the LSAT sub-task, we developed a retrieval system to provide timely, precise and convenient access to a lifelog data archive. The system, as well as our official submissions were designed to maximise recall, in order to support a user to access their life experiences in a real-world lifelogging scenario.

#### 3.1 Data Preprocessing

After analyzing NTCIR-14 lifelog data [3], we divide the data into five main categories: time, location, activities, biometrics and visual concepts.

1. **Time:** For time data, we split the minute based lifelog data into selection of range of hours/minutes/days/day of week for lifelog search engine. Novice/expert user can utilise this information to narrow the scope of searching for a topic in lifelogger’s data. All time data is converted into the UTC time standard.
2. **Location:** For location data, we also utilise timezone information to know the region/country where the lifelogger is visiting. We convert locations into semantic names to help novice/expert user locate the category of place when searching for lifelogger’s moments.
3. **Activity:** The activity data contains two categories: walking and transport.
4. **Biometrics:** The biometrics data that we use in our search engine includes heart rate and calories.
5. **Visual concepts:** We included the visual concepts that were distributed with the dataset. Visual concepts are of three types: place attribute, place category, and visual objects. The place features were extracted using places365-CNN [20]. The visual objects’ categories originate from MSCOCO dataset [11] and are automatically detected using object detection network [17].

### 3.2 Supporting Faceted Search

The interactive retrieval engine implemented a faceted search system in which a user could either enter a textual query in a conventional text box, or select from a range of facets of the metadata to locate items of interest. The faceted search system operated over a range of metadata which are listed in subsection 3.1 which are day of the week, date, time range, user activity (walking/transporting), biometrics data ranges (calories and heart rate), location (location category and name), and visual concepts (place attributes, place categories, and detected objects) in the corresponding order.

When searching using the conventional text box, a user is limited to utilising only visual concepts, activity, and location, and as such, it was a simplified version of a conventional bag-of-words retrieval system. If user desired to utilise all the metadata in searching for relevant items, the faceted query mechanism was required to support this.

The interface, showing the faceted panel (left), the querybox (top) and the result browsing panel (right) is shown in Figure 1. Note the timer on the top of the main panel, which was added to support the LSAT interactive experiment.

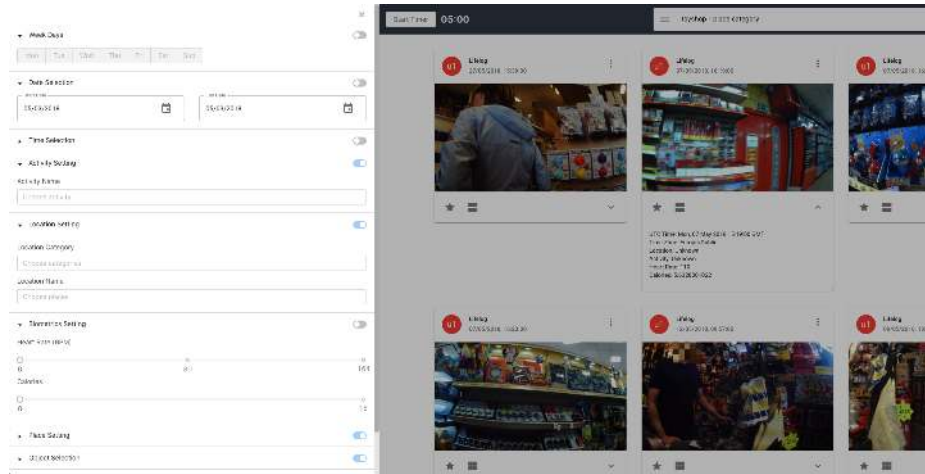
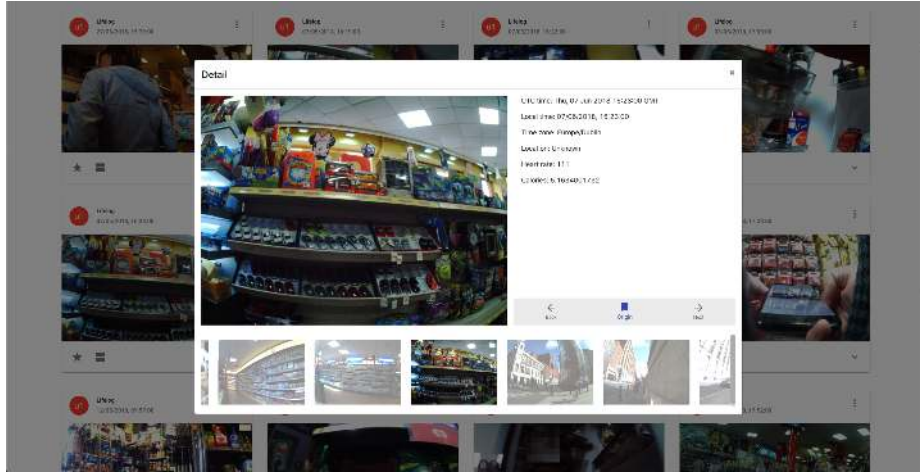


Fig. 1. Query Panel

Upon generating a query, the system generated a list of results (20 per page and 5 pages of results) ranked in temporal order, as shown in 2, using a conventional text ranking algorithm. The unit of retrieval was the image, as was expected for the LSAT task. Each image is given a title, date, a button to select the image and another one to show before & after the current image. Summary metadata from each image could be displayed by selecting the image. If an image was selected as being relevant (the star icon), then it was saved for submission.





**Fig. 3.** Temporal Browsing

was given five minutes to complete each topic. The experimental protocol was as follows. The participant was introduced to the system and given a five minute review of functionality. Following this, the participant was allowed to test the system for a further ten minutes with two sample queries. Once the participant was comfortable with the system and how it operated, the user study began with the user reading a topic and the five minute timer was started once the user was comfortable that they understood the topic. All twelve topics were executed in forward order for users 1 and 2, and in reverse order for users 3 and 4. It would take around 90 minutes per user to conduct the whole experiment. In terms of practical experimental configuration, two users took part in the experiment in parallel (1 and 2, followed by 3 and 4).

## 4.2 Results

The user experiment produced two runs; one combining the submissions of DCU-run1 (users 1, 2), and a second for DCU-run2(users 3, 4). DCU-run1 contained submitted results for 22 of the 24 topics, whereas DCU-run2 contained results for 23 of the 24 topics. For missing topics, it means that the user could not find any relevant images which are suitable to the detailed description of the topic. The total number of retrieved relevant results for DCU-run1 was 556, whereas for DCU-run2, it was 1094. DCU-run2 users found significantly more results than DCU-run1, which highlights a variability in how the teams were formed. Interestingly, users 3 & 4 would have scored the system usability lower than users 1 & 2, although their interaction found over double the number of relevant items.

Considering that we were employing pagination of results at 20 per page, the P@10 metric for DCU-run1 was 0.1917 but for DCU-run2, it was 0.2292. Given

the nature of the experiment, exploring results from a ranked list at higher cut-off points was not valuable due to less similarity to the query’s content. In terms of MAP, DCU-run1 was 0.0724, but for DCU-run2 it was 0.1274, once again significantly higher.

When comparing performance of this system with other participants’ interactive retrieval system in the LSAT sub-task (see table 1), it is apparent that the DCU-Run1 underperformed against other runs in terms of MAP and P@10, with only DCU-Run2 performing better than any competitor. It is our conjecture that this was due to the packing of the result submission with the temporal images, which would have reduced the MAP and P@10 scores. Considering the RelRet (Relevant items Recalled) measure, both runs were only bettered by the HCMUS system [15], which was the overall best performing interactive system. Once again, the submission packing would have increased these RelRet scores. Another factor that could be taken into consideration was the application of a five-minute time limit on each topic. Had this been longer, then the scores would likely have changed.

**Table 1.** Comparing DCU-Run1 & 2 with other Interactive Runs, from [3]

Group ID	Run ID	Approach	MAP	P@10	RelRet
NTU	NTU-Run2	Interactive	0.1108	0.3750	464
NTU	NTU-Run3	Interactive	0.1657	0.6833	407
HCMUS	HCMUS-Run1	Interactive	0.3993	0.7917	1444
<b>DCU</b>	<b>DCU-Run1</b>	<b>Interactive</b>	<b>0.0724</b>	<b>0.1917</b>	<b>556</b>
<b>DCU</b>	<b>DCU-Run2</b>	<b>Interactive</b>	<b>0.1274</b>	<b>0.2292</b>	<b>1094</b>

When considering both automatic and interactive retrieval efforts, the best run that automatic retrieval system achieved was from the NTU group (MAP=0.0632, P@10=0.2375, RelRet=293) [2]. Although their P@10 is slightly higher than our baseline interactive retrieval system, its both RelRet and MAP are not as good as our baseline system. However, the high P@10 score gives promising improvement on automatic retrieval search engine. The overall result of NTCIR-14 for both interactive and automatic retrieval manner could be seen in Table 2.

### 4.3 User Feedback

The inter-run comparisons just presented are not very useful when considering how well a system is liked by users. Clearly users 3 and 4 outperformed users 1 and 2. Using a questionnaire (The User Experience Questionnaire - QEU) [7], we sought to get an initial feedback from users about their experiences with the interactive retrieval engine. All four users filled in the simple 8 part questionnaire, which evaluated the system in terms of pragmatic (realistic-use-case) quality and hedonic (pleasantness) quality, with results shown in table 3. In terms of pragmatic quality, the interface was seen as being slightly more (+0.5 from a maximum of 3.0) supportive than obstructive, slightly more easy (+0.3) than

**Table 2.** Comparing DCU-Run1 & 2 with Automatic Runs, from [3]

Group ID	Run ID	Approach	MAP	P@10	RelRet
NTU	NTU-Run1	Automatic	0.0632	0.2375	293
NTU	NTU-Run2	Interactive	0.1108	0.3750	464
NTU	NTU-Run3	Interactive	0.1657	0.6833	407
HCMUS	HCMUS-Run1	Interactive	0.3993	0.7917	1444
QUIK	QUIK-Run1	Automatic	0.0454	0.1958	232
QUIK	QUIK-Run2	Automatic	0.0454	0.1875	232
<b>DCU</b>	<b>DCU-Run1</b>	<b>Interactive</b>	<b>0.0724</b>	<b>0.1917</b>	<b>556</b>
<b>DCU</b>	<b>DCU-Run2</b>	<b>Interactive</b>	<b>0.1274</b>	<b>0.2292</b>	<b>1094</b>

complicated and slightly more clear (+0.3) than confusing. However users felt that it was slightly more inefficient (-0.3) than efficient. In terms of hedonic quality the interface was considered to be significantly more exciting (+1.3) than boring, significantly more interesting (+2.0) than non-interesting, significantly more inventive (+1.3) than conventional and slightly more leading-edge than usual/conventional.

**Table 3.** Pragmatic quality feedback of DCU’s interactive retrieval engine

	Mean	Variance	Std. Dev.	Negative	Positive	Scale
<b>1</b>	0.5	6.3	2.5	Obstructive	Supportive	Pragmatic Quality
<b>2</b>	0.3	7.6	2.8	Complicated	Easy	
<b>3</b>	-0.3	2.9	1.7	Inefficient	Efficient	
<b>4</b>	0.3	2.9	1.7	Confusing	Clear	
<b>5</b>	1.3	4.3	2.1	Boring	Exciting	Hedonic Quality
<b>6</b>	2.0	1.3	1.2	Not Interesting	Interesting	
<b>7</b>	1.3	2.3	1.5	Conventional	Inventive	
<b>8</b>	0.8	2.3	1.5	Usual	Leading Edge	

Exploring the qualitative findings on a per run basis, DCU-run1 users considered that the system was more supportive, easier to user, more efficient and clearer than DCU-run2 users. In terms of hedonic quality, they also found it more exciting, interesting, inventive and leading edge. However, considering the actual runs, these users were significantly less effective when using the system.

This feedback is reasonable because DCU-run2 users have prior experience of developing application system, which is why they expect the search engine to be more effective, clearer, and less complicated in interacting with our system. In contrast, DCU-run1 users understand how our search engine work after training without any further expectation of user interaction and think that our available functions are enough to retrieve the correct moments.

Through feedback and observation of the users using the retrieval system, we gathered findings that are being used to improve the current system for the LSC’19 (Lifelog Search Challenge) comparative benchmarking exercise. The new



system called LifeSeeker [9] is an evolution of this system that incorporated the following updates:

- Taking measures to reduce the lexical gap (between user queries and the indexed concepts within the system) by expanding the indexed terms to include synonyms. Hence, we enriched the output of the visual and biometric concept detectors using a term-expansion (thesaurus-lookup) approach. For example the concept *seaside* would include the following synonyms; *shore, coast, sands, margin, strand, seaside, shingle, lakeside, water’s edge, lido, foreshore, seashore, plage, littoral, sea*.
- Integrating content-similarity to allow the user to find similar looking content for any given image. Such a feature had been used successfully by participants in LSC’18. For this we utilised the Bag-of-Words model to transform visual features into a vector representation for comparing and returning similar images. Extracting visual features from image was done thanks to the Scale-Invariant-Feature-Transform (SIFT) [13] detector and cosine vector distance as a dissimilarity measure.
- Including a more conventional free-text search element and integrating the filter panel as part of the free-text query mechanism. The free-text ranking engine implemented in the system indexed all textual content associated with any image within the collection. In order to reduce the architectural complexity and latency of the system, we choose to use a standard approach to term weighting [18]. For the purposes of this interactive system, both stemming and stop-words were employed. The maximum number of results returned was 1,000, although in a standard configuration, only 100 were displayed to the user in the interface. The top 1,000 images was necessary for the ranking system to support faceted filtering.

These changes were combined with a slightly revised interface to take in to account the richer metadata and the content similarity functionality, as shown in Figures 4, 5, and 6.

In the interactive search competition at LSC2019, this system performed among the top-ranked teams with an overall score of 68, compared to the vitrivr system [19] which was given a score of 100. Interestingly the system significantly closed the gap to the NTCIR-14 system from HCMUS (which also competed at the LSC in 2019) who scored 72 in the competition. Details of the scoring function employed can be found in the review of the LSC 2018 competition [6].

## 5 Conclusions and Future Work

In this paper, we introduced a first-generation prototype of an interactive retrieval engine for lifelog data, that was run at the NTCIR14-Lifelog3 task and enhanced to be competitive in the second LSC Challenge in 2019. The system was a baseline retrieval system that operated over the provided metadata for the collection. The system was evaluated by four users and findings indicate that the system can be effectively used to locate relevant content. User studies showed




Fig. 4. LifeSeeker interface with free-text search


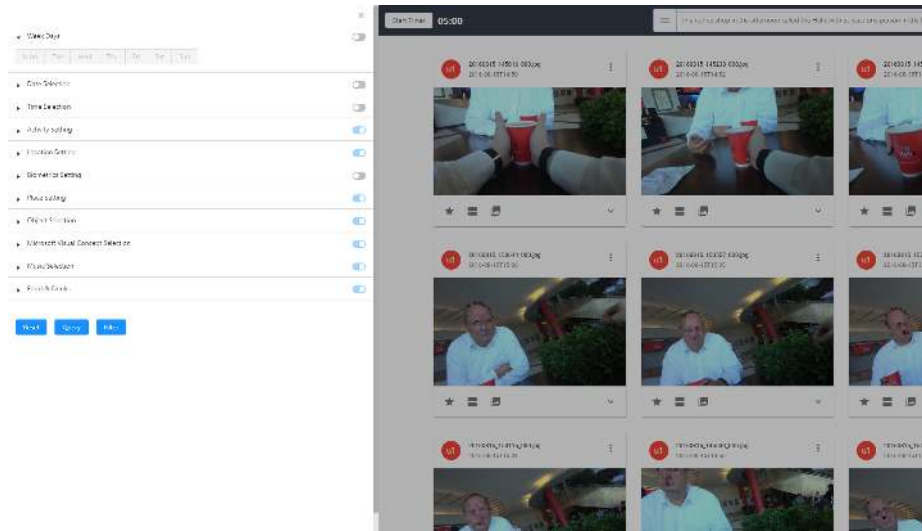


Fig. 5. LifeSeeker interface with visual content similarity function. The items in the grid belong to the image inside the red bounding box

that the users generally liked the system, but both observation and feedback provided a list of proposed enhancements to the system, which have been integrated into a new interactive retrieval system called LifeSeeker [9] which was shown to be among the best performers at the LSC2019.



**Fig. 6.** LifeSeeker interface with additional facet filters

## Acknowledgements

This publication has emanated from research supported in part by research grants from Irish Research Council (IRC) under Grant Number GOIPG/2016/741 and Science Foundation Ireland under grant numbers SFI/12/RC/2289 and 13/RC/2106. We wish to also thank the four users who used our retrieval system to generate the official runs.

## References

1. Duane, A., Gurrin, C., Hürst, W.: Virtual reality lifelog explorer: Lifelog search challenge at acm icmr 2018. In: Proceedings of the 2018 ACM Workshop on The Lifelog Search Challenge. pp. 20–23. LSC '18, ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3210539.3210544>, <http://doi.acm.org/10.1145/3210539.3210544>
2. Fu, M.H., Chia-Chun, C., Huang, G.H., Chen, H.H.: Introducing external textual knowledge for lifelog retrieval and annotation. In: The Fourteenth NTCIR conference (NTCIR-14) (2019)
3. Gurrin, C., Joho, H., Hopfgartner, F., Dang-Nguyen, D.T., Zhou, L., Ninh, V.T., Le, T.K., Albatat, R., Healy, G.: Overview of NTCIR-14 lifelog task. In: Proceedings of the 14th NTCIR Conference on Evaluation of Information Access Technologies, National Center of Sciences, Tokyo, Japan, June 10-13, 2019 (2019)
4. Gurrin, C., Joho, H., Hopfgartner, F., Zhou, L., Albatat, R.: Overview of NTCIR-12 lifelog task. In: Proceedings of the 12th NTCIR Conference on Evaluation of Information Access Technologies, National Center of Sciences, Tokyo, Japan, June 7-10, 2016 (2016)

5. Gurrin, C., Schoeffmann, K., Joho, H., Münzer, B., Albatal, R., Hopfgartner, F., Zhou, L., Dang-Nguyen, D.T.: A Test Collection for Interactive Lifelog Retrieval. In: MultiMedia Modeling - 25th International Conference, {MMM} 2019, Thessaloniki, Greece, January 8-11, 2019, Proceedings, Part {I}. Lecture Notes in Computer Science, vol. 11295. Springer (2019)
6. Gurrin, C., Schoeffmann, K., Joho, H., Zhou, L., Duane, A., Leibetseder, A., Riegler, M., Piras, L.: Comparing Approaches to Interactive Lifelog Search at the Lifelog Search Challenge ( LSC2018 ). *ITE Transactions on Media Technology and Applications* **7**(2), 46–59 (2019)
7. Laugwitz, B., Held, T., Schrepp, M.: Construction and evaluation of a user experience questionnaire. In: Holzinger, A. (ed.) *HCI and Usability for Education and Work*. pp. 63–76. Springer Berlin Heidelberg, Berlin, Heidelberg (2008)
8. Le, N.K., Nguyen, D.H., Hoang, T.H., Nguyen, T.A., Truong, T.D., Dinh, D.T., Luong, Q.A., Vo-Ho, V.K., Nguyen, V.T., Tran, M.T.: Smart lifelog retrieval system with habit-based concepts and moment visualization. In: Proceedings of the ACM Workshop on Lifelog Search Challenge. pp. 1–6. LSC '19, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3326460.3329155>, <http://doi.acm.org/10.1145/3326460.3329155>
9. Le, T.K., Ninh, V.T., Dang-Nguyen, D.T., Tran, M.T., Zhou, L., Redondo, P., Smyth, S., Gurrin, C.: Lifeseeker - interactive lifelog search engine at lsc 2019. In: Proceedings of the 2019 ACM Workshop on The Lifelog Search Challenge. LSC '19, ACM, New York, NY, USA (2019)
10. Leibetseder, A., Münzer, B., Primus, M.J., Kletz, S., Schoeffmann, K., Berns, F., Beecks, C.: lifexplore at the lifelog search challenge 2019. In: Proceedings of the ACM Workshop on Lifelog Search Challenge. pp. 13–17. LSC '19, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3326460.3329157>, <http://doi.acm.org/10.1145/3326460.3329157>
11. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context pp. 740–755 (2014)
12. Lokoč, J., Souček, T., Kovalčík, G.: Using an interactive video retrieval tool for lifelog data. In: Proceedings of the 2018 ACM Workshop on The Lifelog Search Challenge. pp. 15–19. LSC '18, ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3210539.3210543>, <http://doi.acm.org/10.1145/3210539.3210543>
13. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision* **60**(2), 91–110 (Nov 2004). <https://doi.org/10.1023/B:VISI.0000029664.99615.94>, <http://doi.org/10.1023/B:VISI.0000029664.99615.94>
14. Münzer, B., Leibetseder, A., Kletz, S., Primus, M.J., Schoeffmann, K.: lifexplore at the lifelog search challenge 2018. In: Proceedings of the 2018 ACM Workshop on The Lifelog Search Challenge. pp. 3–8. LSC '18, ACM, New York, NY, USA (2018). <https://doi.org/10.1145/3210539.3210541>, <http://doi.acm.org/10.1145/3210539.3210541>
15. Nguyen, I.V.K., Shrestha, P., Zhang, M., Liu, Y., Ma, S.: Thuir at the ntcir-14 lifelog-3 task: How does lifelog help the users status recognition. In: The Fourteenth NTCIR conference (NTCIR-14) (2019)
16. de Oliveira Barra, G., Ayala, A.C., Bolaños, M., Dimiccoli, M., Giró i Nieto, X., Radeva, P.: Lemore: A lifelog engine for moments retrieval at the ntcir-lifelog LSAT task. In: Proceedings of the 12th NTCIR Conference on Evaluation of Information Access Technologies, National Center of Sciences, Tokyo, Japan, June 7-10, 2016 (2016)

17. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1. pp. 91–99. NIPS'15, MIT Press, Cambridge, MA, USA (2015), <http://dl.acm.org/citation.cfm?id=2969239.2969250>
18. Robertson, S.E., Jones, K.S.: Simple, proven approaches to text retrieval. Tech. rep. (1997)
19. Rossetto, L., Gasser, R., Heller, S., Amiri Parian, M., Schuldt, H.: Retrieval of structured and unstructured data with vitrivr. In: Proceedings of the ACM Workshop on Lifelog Search Challenge. pp. 27–31. LSC '19, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3326460.3329160>, <http://doi.acm.org/10.1145/3326460.3329160>
20. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2017)