

Lifelogging As An Extreme Form of Personal Information Management - What Lessons To Learn

Ly-Duyen Tran¹ Cathal Gurrin^{1,2} Alan F. Smeaton^{1,3*}

¹School of Computing

²ADAPT Centre

³Insight Centre for Data Analytics

Dublin City University, Glasnevin, Dublin 9
Ireland.

Abstract

Personal data includes the digital footprints that we leave behind as part of our everyday activities, both online and offline in the real world. It includes data we collect ourselves, such as from wearables, as well as the data collected by others about our online behaviour and activities. Sometimes we are able to use the personal data we ourselves collect, in order to examine some parts of our lives but for the most part, our personal data is leveraged by third parties including internet companies, for services like targeted advertising and recommendations. Lifelogging is a form of extreme personal data gathering and in this article we present an overview of the tools used to manage access to lifelogs as demonstrated at the most recent of the annual Lifelog Search Challenge benchmarking workshops. Here, experimental systems are showcased in live, real time information seeking tasks by real users. This overview of these systems' capabilities show the range of possibilities for accessing our own personal data which may, in time, become more easily available as consumer-level services.

1 What is Personal Information Management ?

Personal information or personal data is the digital form of the footprints of the everyday things that we do in our daily lives. It includes data that we can generate ourselves for example from wearable devices, or information that we ourselves create directly like our emails. In the case of the former, data from wearable devices can include physiological indicators like heart rate, respiration rate, heart rate variability or galvanic skin response levels, which are indicators of stress. It can also include our movement data from accelerometers on as well as location data from GPS or other location-tracking technologies. When such raw data is analysed then we can infer things like step counts, distances walked, run or biked, metrics for our sleep like duration and sleep quality and other health and wellness indicators. The purpose of capturing such personal data about ourselves is usually self-monitoring and self-tracking, typically for our own health awareness.

While self-tracking is now a consumer-level rather than a specialist activity with hundreds of thousands of apps for uses like counting calories or fitness tracking [28] much of what we call our personal data is actually recorded by others and not directly by ourselves. This includes the logs and records of our online activities like logs of our web browsing, social media interactions, media consumption and our communications and interactions with others. Even though this data has been gathered and is being used by others, it is personal data that belongs to us. Personal data gathered by third parties has value because of the insights it can reveal about us and especially

*Contact: Alan.Smeaton@DCU.ie

when it is aggregated and combined with the personal data of others [6]. This form of personal data is exploited almost exclusively by others, albeit supposedly for our benefit through targeted advertising and recommendations, and not directly by us.

Personal information management refers to tools and platforms that allow us to control our personal data by allowing us to gather, store, update, analyse, interpret and sometimes to share it, no matter who has gathered it. It is awkward or at best inconvenient for us to access or even to use our own personal data for our own personal purposes whether it is data we have gathered, or has been gathered by third parties. While the log files of our interactions with search, browsing, social media, communications and media consumption services, were originally intended for helping to debug and refine those systems and while we can request such logs from internet services it is not made easy for us to do so. Even for our use of systems in our own enterprises and organisations, like access logs to virtual learning environments in Colleges and Universities, this typically requires downloading CSV files of personal data and then processing them ourselves, which is off-putting to most. Access to our personal data from wearable devices is also typically supported through offering CSV or equivalent files as downloads.

It is clear that our personal data can be a rich source of insights into us, into our well-being, our habits and behaviours and most importantly, into changes into those behaviours. Unfortunately we are able to query and to interrogate the individual and unconnected repositories for our personal data only on a per-device or per-source basis whereas the real benefits of such interrogations would be when the whole of our personal data, integrated across sources, could be analysed. This poses the question of whether there is any work done that brings these sources together.

In the later part of this article we will look at some very limited attempts to pool together our personal data from multiple sources but before that, in the next section, we will examine the current best-practice in the area of self-generated recording of everyday activities, the area known as lifelogging.

2 Lifelogging: Extreme Personal Information Management

In this section, we take a closer look as an extreme form of personal information gathering, which is called lifelogging [14]. Lifelogging involves the comprehensive and continuous collection of data about one's daily activities and experiences, gathered from a wide array of data sources. Wearable devices, including smartwatches and fitness trackers, can be used to capture health metrics, such as heart rate, sleep quality, and activity levels. GPS sensors can be used to track one's movements and location. The most prominent data source for lifelogging is the use of wearable cameras to capture point-of-view images and videos of one's daily activities. These data sources provide a holistic representation of one's life and can be used for a variety of applications, such as health monitoring, memory prosthetic, and behavioural analysis.

The applications of lifelogging are widespread and have been explored in many domains, such as supporting human memory recollection [3, 5, 15], supporting large-scale epidemiological studies in healthcare [35], monitoring the lifestyle of individuals [29, 48], behaviour analytics [12], and diet/obesity analytics [51]. However, the vast volume and immense complexity of lifelog archives presents a challenge for users to navigate and analyse these archives in order to identify relevant information. As such, there is a growing interest in developing lifelog retrieval systems that effectively leverage lifelog data to meet the diverse needs of users.

Retrieval systems for lifelog data have been a popular research topic for several years. Managing the large amount of data collected by lifeloggers is crucial in order to maximise benefits from the data. In the MyLifeBits system, introduced more than 20 years ago by Gordon Bell [4], then state-of-the-art techniques from database search and traditional information retrieval (IR) were employed to index and provide access to lifelog data through a desktop interface. However, as the volume of lifelog data increases, the need for more efficient and effective retrieval systems becomes more apparent and the limitations of MyLifeBits soon became apparent.

The first retrieval system designed for large archives of lifelog data was proposed by Doherty *et al.* [11],

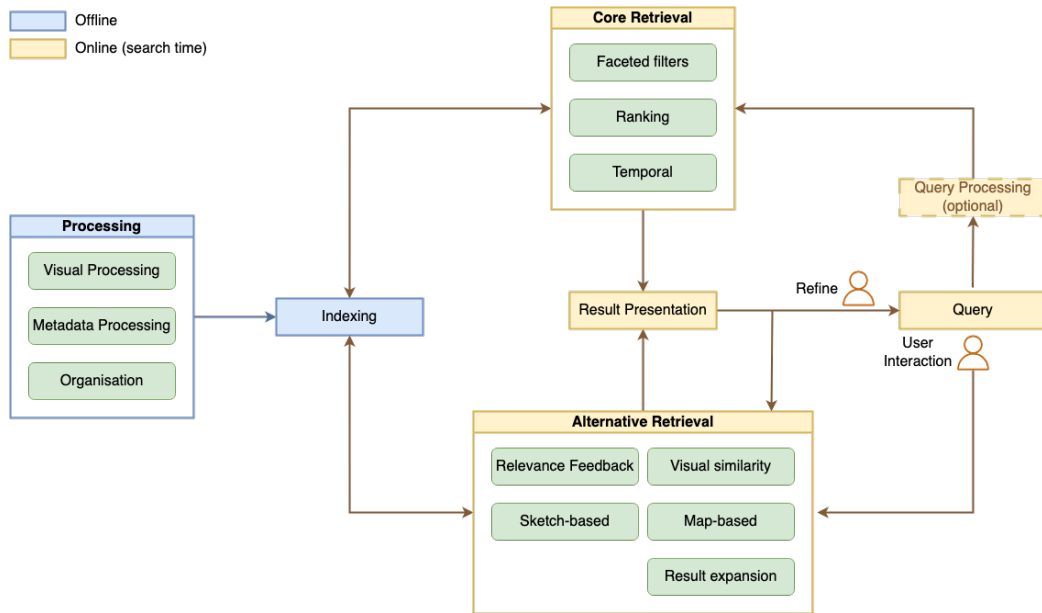


Figure 1: General pipeline of lifelog retrieval systems in the annual Lifelog Search Challenge (LSC).

moving the time/date *browsing* approach of lifelog systems at the time to a *search* approach. The system employed event segmentation, event annotation and multi-axes search, which are the ‘who’, ‘what’, ‘when’, and ‘where’ axes of retrieval. However, without a large user base, it was difficult to define search use-cases in order to evaluate and improve lifelog retrieval systems.

In order to explore the possibilities and limitations of question-answering on the unique characteristics of timeline information such as that gathered in lifelogs, the authors of [39] have created and released TimelineQA. This is a generator that produces synthetic lifelogs for imaginary people with different personae including age, gender, education and family status. It uses a fine-tuned LLM to produce the actual content on which the effectiveness of question-answering can be investigated. The synthetic lifelogs are composed of timelines of ordinary life events each of which have start times, end times and possibly locations, and they may also have multimodal features like photos. The TimelineQA benchmark is successful in that it is allowing a new form of interaction with lifelogs – question-answering – to be explored in a systematic and repeatable way.

Several other benchmarking activities for lifelog retrieval systems have been organised which evaluate and compare the retrieval effectiveness of participating systems on real rather than synthetic lifelogs. The notable and best example of this is the annual Lifelog Search Challenge (LSC) [42]. The LSC provides a platform for evaluating state-of-the-art systems for managing lifelogs and attracts a large number of participants internationally. In the LSC, different systems compete with each other in a live/virtual environment, where participants are given a set of information needs and queries and a limited time to find the relevant and matching lifelog moments, in real time. Generally, most systems participating in the LSC follow the flow of data processing, indexing, and then retrieval, as shown in Figure 1. Interactivity and user experience are important aspects of the lifelog retrieval systems, as their affordances are known to directly influence the effectiveness of the retrieval process [26]. We now discuss some of the techniques and components involved in lifelog retrieval with example systems from the LSC.

2.1 Processing

Lifelog data is typically collected from sources including wearable cameras, wearable sensors, and mobile devices. The data is then processed to extract relevant information before being indexed and stored in a database. The

processing stage is crucial as it determines the quality of the data and the effectiveness of the retrieval process.

2.1.1 Visual Processing

The most challenging aspect of processing is visual data, where different techniques are employed. We categorise these techniques into three groups: low-level features, concepts, captions, and embeddings.

Low-level features, such as colours, textures, and shapes, are commonly used in lifelog systems. These are classic features that have been used in the field of image processing for decades and are still used in lifelog systems. Their most popular application is the removal of low-quality images as seen in systems such as MyScéal [43] (or its variant, E-MyScéal) and Memento [1]. It is believed that low-quality images are less likely to be relevant to users, and thus, removing them can improve retrieval performance.

Due to advances in computer vision algorithms, many lifelog systems utilise deep convolutional neural networks (CNNs) such as AlexNet [23], VGG [36], GoogLeNet [38], and ResNet [16], trained on various dataset such as ImageNet [8], OpenImages [24], Places365 [49], and Visual Genome [22]. These models are used to extract semantic lifelog ‘concepts’, which are then indexed and used for retrieval. Optical character recognition (OCR) is also used to extract text from images, which proved to be useful in scenarios like identifying brand names and street names [42] during the real time search. This approach, often referred to as concept-based retrieval, is widely adopted due to its ease of implementation and the availability of pre-trained models.

Captioning, the generation of textual descriptions for images, is less commonly used in lifelog systems due to its inherent challenges. Generated captions are not always accurate, which limits its popularity. The most notable example of captioning in lifelog systems is vitivr [37], which supports text-based search using generated captions.

Embedding-based retrieval is an approach that has gained popularity in recent years. This involves mapping images into a vector space, where the similarity between images can be measured using cosine similarity, for example. This approach is widely used in cross-modal retrieval, where images are mapped into the same vector space as text queries. The most notable example of this approach is CLIP [32], which is a pre-trained model that maps images and text into a shared vector space. This approach replaces the need for concept extraction and captioning, as the similarity between images and text can be measured directly. Therefore, many lifelog systems have replaced their concept-based retrieval with embedding-based retrieval, as demonstrated by E-MyScéal [43], by adding embeddings to the existing concept-based retrieval and thus the name E-MyScéal. Other systems such as Memento [1] and LifeSeeker [30] have also adopted this approach. This *embedding-based retrieval* approach allows a more user-friendly experience by allowing users to search for images using natural language queries and is shown to significantly improve retrieval performance

2.1.2 Metadata Processing

While metadata processing in lifelog systems is minimal due to standardised CSV formats provided by the LSC challenge organisers, it still plays a vital role in enhancing retrieval performance. The two most common metadata used in lifelog systems are timestamps and GPS coordinates. These help at retrieval time as users often remember aspects of the date or time or of the approximate location, of their personal data they are trying to find. Time information such as the time of the day, day of the week, month, and year are extracted from timestamps and used for filtering and scoring, as demonstrated by MyScéal [43] and LifeSeeker [30]. Clustering GPS coordinates for stay point detection, and inferring semantic locations are common techniques used to enhance the location metadata. For example, MyEachtra [41] employs a visual-aware stay point detection algorithm [44] to detect stay points and utilises crowdsourced location data from FourSquare¹ to infer semantic locations. Meanwhile, LifeSeeker [30] enhances the location metadata by manually labelling images with ‘areas’ within a location (such

¹<https://developer.foursquare.com/docs/api-reference/>

as kitchen, living room, and bedroom) to create a more fine-grained location hierarchy. This approach provides more ‘concepts’ for users to filter and search for.

Biometric sensor data, surprisingly, is rarely used in lifelog systems due to noisy data and the lack of standardised formats. However, some systems have demonstrated the potential of using biometric data for retrieval by binning the data into categories [2, 31, 45]. Another type of metadata that was used in lifelog systems is music listening data. LifeSeeker uses music listening data to infer a user’s mood, which is then used to filter and score the search results [30]. However, due to the lack of LSC queries with the focus on these types of metadata, their effectiveness is yet to be evaluated.

2.1.3 Organisation

Due to the temporal nature of lifelog data, it is straightforward to organise lifelog data chronologically [50]. Some of the LSC participating systems use a more sophisticated approach to organise the data, such as by using segmentation. Most systems that include segmentation use a simple approach by comparing the visual similarity between two consecutive images [43]. MyEachtra [41] clusters the GPS coordinates to infer the semantic locations and then organise the data by the semantic locations.

LifeGraph [33] offers a novel approach and organises the lifelog data by using a graph structure, with images as the centrepiece of the schema. The graph is constructed by linking the images with the metadata and detected concepts, then is extended with Wikidata [47] and COEL (Classification of Everyday Living) [7]. However, the authors later acknowledged that COEL played an insignificant role in the query expansion and that Wikidata was a more useful source of concepts.

2.2 Search and Navigation

The core aspect distinguishing various lifelog retrieval systems is their retrieval mechanism. This significantly influences how data is processed, indexed, and presented.

Filters are a widely employed technique that has proven to be effective in lifelog systems. Faceted filters can be applied to various modalities and are often categorised by attributes such as time of day, day of the week, month, year, location, number of people, biometrics, and lifelog concepts as these are the aspects of personal data that a user can recall in their information seeking [21]. Fixed phrases are typically used for filtering values, presented through drop-down lists, checkboxes, or sliders. Some systems, like MyScéal [43], allow users to enter free text, with autocomplete suggestions for available lifelog concepts.

Query expansion is a technique used to assist users in formulating queries. It typically involves concept suggestion using various sources like WordNet, ConceptNet, and Thesaurus.com, and employing models like BERT [9] for concept similarity. MyScéal [43] introduced a free-text query form with filter value extraction and query expansion, a method later adopted by LifeSeeker [30].

A ranking mechanism is essential for any retrieval system in order to sort outputs before presentation. Lifelog concepts can be used to score images based on the number of matched concepts [50]. Extending TF-IDF term weighting using concepts extracted from computer vision models is a popular approach, as demonstrated by MyScéal [43] and by LifeSeeker [30]. TF-IDF is a common technique in information retrieval to score the relevance of a document to a query. It is calculated based on the term frequency (TF) and inverse document frequency (IDF) of the query terms. TF scores are not as useful as they are in the field of IR since the terms (concepts) are oftentimes not repeated in a document (image). Therefore, the TF scores are often replaced by the confidence scores of the concepts extracted from various computer vision models [10]. The area of the object (or its bounding box) can also be exploited as in aTFIDF, proposed by MyScéal [43].

As previously mentioned, cross-modal embedding models have offered a new way of measuring the similarity between images and text in lifelog retrieval systems. Cosine similarity between the search query and the images is directly used to rank the result. Some optimisation methods such as FAISS [19] or KNN search [43] can be used

to speed up the search. Although there is some effort of fine-tuning the embedding models on lifelog models [40], large-scale pre-trained models are more robust and are often used directly or in a weighted ensemble in many systems, with Memento [1] being the most notable example.

In addition, support for temporal search is available in many LSC systems, enabling users to combine temporally related queries. For example, a user can search for ‘eating apple before watching TV’. Multiple temporally ordered queries can be executed conditionally, with results re-ranked accordingly based on the scores of each query.

2.3 User Interactions

Relevance feedback mechanisms are utilised in several lifelog systems to enhance retrieval performance. This technique, commonly employed in Information Retrieval, enables users to provide feedback on search results. Users can label images as relevant or irrelevant, prompting the system to refine the search iteratively. For example, Exquisitor [20] incorporates relevance feedback, where a user’s feedback influences the training of classifiers to retrieve a more relevant set of images. Similar mechanisms are applied in systems like SOMHunter [27], where a user selects relevant images, or in Exclusion of concepts, allowing users to specify concepts they wish to exclude from the results.

Visual similarity plays a significant role in lifelog retrieval systems, often used in conjunction with other retrieval methods. It can be employed to arrange result lists as seen in SOMHunter [27] and lifeExplore [34]; or to offer an alternative means for users to explore lifelog archives as seen in MyScéal [43] and LifeSeeker [30]. Visual similarity can be computed based on various factors, including low-level features such as colour histograms and SIFT, content-based features, or cross-modal embedding techniques.

Sketch-based search is a compelling method implemented in some lifelog systems originating from the video search community. Systems like lifeExplore [34] and vitivr [37] allow users to draw a sketch of the image they seek. The system then returns visually similar images to the sketch. This method proves useful when users cannot describe an image they wish to locate in words but is less popular due to users’ limited ability to convey a target image accurately.

Map-based search is another frequently employed technique in lifelog systems. Users can draw a rectangle on a map to narrow down the search space, focusing only on moments that occurred within that area. Systems like lifeExplore [34], MyScéal [43], and vitivr [37] offer map-based search features, often utilising libraries like Leaflet.

2.4 Result Presentation Techniques

Result presentation in lifelog retrieval systems has evolved beyond conventional grid views. Some systems have explored alternative ways of presenting results to enhance user experience. For instance, lifeExplore [34] and SOMHunter [27] utilise Self-Organising Maps (SOMs) to arrange images in a 2D map, clustering them based on visual similarity. An autopilot navigation mode, introduced by lifeExplore in LSC’20, ensures every image in the target area is ‘visited’ once.

To mitigate visual clutter caused by similar images in lifelog data, event clustering is employed. Systems like MyScéal [43] perform offline event segmentation and display the highest-ranked image from each event. In contrast, LifeSeeker [30] dynamically clusters images during retrieval based on visual similarity and temporal proximity, presenting the top three images from each cluster.

Addressing the temporal aspect of lifelogs, MyScéal [43] proposes presenting results in triplets, showing the immediate previous and next events alongside the target event to provide more context. However, this approach is selectively used, specifically when temporal queries are specified, as adapted by E-MyScéal in LSC in 2022.

Memento [1] offers a unique visualisation of results through distribution charts, allowing users to modify filters directly by interacting with the charts. Transitional graph-based visualisation, proposed by LifeSeeker [30],

shows location transitions between images in the result list, which is particularly useful for transportation-related queries.

The ability to browse a lifelog archive chronologically, often referred to as the ‘timeline view’, is crucial for users to explore their lifelog data. Many systems support this feature with varying levels of granularity and designs. Exquisitor [20], for example, designed a temporal context view resembling a video player with lifelog images as thumbnails. Users can play the video to view images chronologically and navigate to specific times by selecting thumbnails below the video. Adjusting the spacing between thumbnails is also possible, which can be implemented using different methods such as scaling factors, step sliders, or hierarchical levels, as seen in systems like LifeSeeker [30] and MyScéal [43].

2.5 Summary

In summary, lifelogging, as an extreme form of personal information management, offers a unique opportunity to gather and manage personal information. The annual lifelog search challenge event (LSC) serves as a remarkable example of how lifelog data can be harnessed and managed effectively. The techniques and components involved in the LSC represent the cutting edge of lifelog management, providing valuable insights into the future of personal data management.

3 Searching Personal Information

In the previous section we have seen the potential for analysing and searching lifelogs as demonstrated in the annual Lifelog Search Challenge. All the systems participating in this interactive live benchmarking event are highly engineered and specialist and not designed for consumer-grade use, yet.

In terms of our own personal data, we know that we can query and interrogate the individual, silo-ed and unconnected repositories which have generated that personal data which has come from individual sources, but is there anything in the personal information landscape that brings these together in some way to create an holistic overview ? A recent review of some personal data stores which can provide access to data from our online activities can be found at [13]. Some example systems included in that review include:

- **Digi.me:** A platform that lets users collect their personal data from online sources including social media, health, finance, and music, and to store it securely on their own devices or on the Digi.me cloud services. Users can then share their data with third parties on their own terms and get personalised insights and benefits from their data [17];
- **Solid:** A project that aims to give users full ownership and control of their data by creating a decentralised web where they can store their data in personal online data stores called pods and link them to applications that respect their privacy and preferences [18].;
- **Meeco:** A platform that enables users to create a digital identity wallet where they can store and manage their personal data such as identity documents, credentials, preferences, and consent. Users can then use their wallet to access online services and share their data securely and selectively [13].

While these examples of personal data stores may be useful for some use cases, they require the user her/himself to carry out the analysis of behaviour changes or patterns, and they are data stores rather than analysis platforms. Thus they are not for non-technical users. As an alternative to managing and analysing our own personal data we could invoke a third party to do the analysis, independently of the large internet companies. A platform to do this has been developed and demonstrated by Tuovinen [46] but that work is still early stage and not yet scaled up.

In the previous section we summarised the state-of-the-art in interactive lifelog search systems and showed that these systems have interactive question-answering capabilities which can lead to conversational retrieval

around the topic of our own personal data. Outside the scope of the Lifelog search Challenge, the tools available to us for querying our own personal data do not yet exist for widespread use and those that are available do not have any forms of aggregation, or recommendation, or behavioural analysis, or self-reflection.

Within many areas of computing, the last year has seen a huge upsurge of interest in the use of large language models (LLMs) for applications in media creation, health and medicine, language processing, teaching and education, social interactions, science, and almost any domain where technology is used. While ChatGPT gets all the media coverage and popular focus, it is the fine-tuning of LLMs on narrow applications which can be restricted in their scope, that holds the greatest long-term potential. The use of the Retrieval Augmented Generation (RAG) [25] which allows an LLM to be fine-tuned on a limited training set could be applied to fine-tuning of an individual's personal data. This would allow the resulting system to leverage the advantages of LLMs including the capability for contextual understanding allowing a better understanding of a user's information need and search intent as well as the ability to simulate dialogue. Through careful prompt engineering it can also be used to unify across different sources of personal data and provide unified search, browsing and navigation across sources. The first example we see of this is when in September 2023 Google announced that BARD AI, their conversational retrieval tool which uses Google's LLM, can integrate with a user's emails, documents and other files as well as internet sources, to give personalised answers to questions. It is thus an obvious next step that conversational retrieval supported by large language models, will be used by the major internet companies as they way for us to access our personal data and that will certainly raise issues of data protection and privacy.

4 Conclusions

Personal information management involves the collection, storage, and control of digital data that represents our daily activities, including data generated by wearable devices, emails, and online interactions. This data can encompass various aspects of our lives, from physiological markers and indicators to our online behaviours. The purpose of gathering this personal data is primarily self-monitoring and self-tracking for health awareness. However, a significant portion of our personal data is collected and utilised by third parties, for targeted advertising and recommendations.

Personal information management tools and platforms should empower individuals to control our personal data, enabling us to gather, store, update, analyze, interpret, and sometimes share this data, regardless of its source. The development of such tools is essential because accessing and using personal data, even when it is our own data, is currently inconvenient and challenging due to the dispersed nature of our data across multiple sources. The paper presents an overview of an extreme form of personal data gathering called lifelogging and in particular how the annual lifelog search challenge highlights how innovative and useful searching through our lifelogs, our personal data, can be.

The article also highlights the potential of fine-tuning large language models (LLMs) for personalised data management, using techniques like Retrieval Augmented Generation (RAG). This approach could enable better contextual understanding, unified searching across various personal data sources, and the potential for conversational retrieval, thus revolutionising how individuals could access and manage their personal data. With the rise of large language models and conversational retrieval, major internet companies are likely to play a significant role in providing access to our personal data, raising important questions regarding data protection and privacy in the future.

References

- [1] N. Alam, Y. Graham, and C. Gurrin. Memento 3.0: An enhanced lifelog search engine for LSC'23. In Proceedings of the 6th Annual ACM Lifelog Search Challenge, pages 41–46. Association for Computing Machinery, New York, NY, USA, 2023.

- [2] A. Alsina, X. Giró, and C. Gurrin. An interactive lifelog search engine for LSC 2018. In Proceedings of the 2018 ACM Workshop on The Lifelog Search Challenge, pages 30–32, New York, NY, USA, 2018. Association for Computing Machinery.
- [3] P. J. Barnard, F. C. Murphy, M. T. Carthery-Goulart, C. Ramponi, and L. Clare. Exploring the basis and boundary conditions of SenseCam-facilitated recollection. Memory, 19(7):758–767, 2011.
- [4] G. Bell. A personal digital store. Commun. ACM, 44(1):86–91, Jan 2001.
- [5] E. Berry, N. Kapur, L. Williams, S. Hodges, P. Watson, G. Smyth, J. Srinivasan, R. Smith, B. Wilson, and K. Wood. The use of a wearable camera, SenseCam, as a pictorial diary to improve autobiographical memory in a patient with limbic encephalitis: A preliminary report. Neuropsychological Rehabilitation, 17(4-5):582–601, 2007.
- [6] K. Birch, D. Cochrane, and C. Ward. Data as asset? the measurement, governance, and valuation of digital personal data by big tech. Big Data & Society, 8(1):20539517211017308, 2021.
- [7] P. Bruton, J. Langford, M. Reed, and D. Snelling. Classification of Everyday Living Version 1.0, 2019.
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 248–255. IEEE, 2009.
- [9] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805, 2018.
- [10] M. Dogariu and B. Ionescu. Multimedia Lab @ ImageCLEF 2018 Lifelog Moment Retrieval Task. In CLEF (Working Notes), page 13, 2018.
- [11] A. R. Doherty, K. Pauly-Takacs, N. Caprani, C. Gurrin, C. J. Moulin, N. E. O’Connor, and A. F. Smeaton. Experiences of aiding autobiographical memory using the SenseCam. Human–Computer Interaction, 27(1-2):151–174, 2012.
- [12] B. Everson, K. A. Mackintosh, M. A. McNarry, C. Todd, and G. Stratton. Can wearable cameras be used to validate school-aged children’s lifestyle behaviours? Children, 6(2):20, 2019.
- [13] K. U. Fallatah, M. Barhamgi, and C. Perera. Personal Data Stores (PDS): A Review. Sensors, 23(3), 2023.
- [14] C. Gurrin, A. F. Smeaton, and A. R. Doherty. Lifelogging: Personal big data. Foundations and Trends in Information Retrieval, 8(1):1–125, 2014.
- [15] M. Harvey, M. Langheinrich, and G. Ward. Remembering through lifelogging: A survey of human memory augmentation. Pervasive and Mobile Computing, 27:14–26, 2016.
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 770–778, 2016.
- [17] H. Janssen, J. Cobbe, and J. Singh. Personal information management systems: a user-centric privacy utopia? Published in Internet Policy Review (18 December 2020), 9(4):1–25, 2020.
- [18] M. Jesús-Azabal, J. Berrocal, S. Laso, J. M. Murillo, and J. Garcia-Alonso. SOLID and PeaaS: Your Phone as a Store for Personal Data. In Current Trends in Web Engineering: ICWE 2020 International Workshops, KDWEB, Sem4Tra, and WoT4H, Helsinki, Finland, June 9–12, 2020, Revised Selected Papers 20, pages 5–10. Springer, 2020.

- [19] J. Johnson, M. Douze, and H. Jégou. Billion-scale similarity search with GPUs. IEEE Transactions on Big Data, 7(3):535–547, 2019.
- [20] O. S. Khan, A. Duane, B. Þ. Jónsson, J. Zahálka, S. Rudinac, and M. Worring. Exquisitor at the lifelog search challenge 2021: Relationships between semantic classifiers. In Proceedings of the 4th Annual on Lifelog Search Challenge, pages 3–6. Association for Computing Machinery, New York, NY, USA, 2021.
- [21] L. M. Koesten, E. Kacprzak, J. F. A. Tennison, and E. Simperl. The trials and tribulations of working with structured data: A study on information seeking behaviour. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI ’17, page 1277–1289, New York, NY, USA, 2017. Association for Computing Machinery.
- [22] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. S. Bernstein, and F.-F. Li. Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations. arXiv:1602.07332 [cs], Feb. 2016. arXiv: 1602.07332.
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems, 25, 2012.
- [24] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Mallocci, A. Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. International Journal of Computer Vision, 128(7):1956–1981, 2020.
- [25] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33:9459–9474, 2020.
- [26] J. Liu, M. Mitsui, N. J. Belkin, and C. Shah. Task, information seeking intentions, and user behavior: Toward a multi-level understanding of web search. In Proceedings of the 2019 Conference on Human Information Interaction and Retrieval, pages 123–132, 2019.
- [27] J. Lokoč, F. Mejzlik, P. Veselý, and T. Souček. Enhanced SOMHunter for known-item search in lifelog data. In Proceedings of the 4th Annual on Lifelog Search Challenge, pages 71–73. Association for Computing Machinery, New York, NY, USA, 2021.
- [28] D. Lupton. Self-tracking, health and medicine, 2017.
- [29] T.-H.-C. Nguyen, J.-C. Nebel, and F. Florez-Revuelta. Recognition of activities of daily living with egocentric vision: A review. Sensors, 16(1):72, 2016.
- [30] T.-N. Nguyen, T.-K. Le, V.-T. Ninh, C. Gurrin, M.-T. Tran, T. B. Nguyen, G. Healy, A. Caputo, and S. Smyth. E-LifeSeeker: An interactive lifelog search engine for lsc’23. In Proceedings of the 6th Annual ACM Lifelog Search Challenge, pages 13–17. Association for Computing Machinery, New York, NY, USA, 2023.
- [31] I. Nguyen Van Khan, P. Shrestha, M. Zhang, Y. Liu, and S. Ma. A two-level lifelog search engine at the LSC 2019. In Proceedings of the ACM Workshop on Lifelog Search Challenge, pages 19–23, 2019.
- [32] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In International Conference on Machine Learning, pages 8748–8763. PMLR, 2021.

- [33] L. Rossetto, O. Inel, S. Lange, F. Ruosch, R. Wang, and A. Bernstein. Multi-Mode Clustering for Graph-Based Lifelog Retrieval. In Proceedings of the 6th Annual ACM Lifelog Search Challenge, pages 36–40. Association for Computing Machinery, New York, NY, USA, 2023.
- [34] K. Schoeffmann. LifeXplore at the Lifelog Search Challenge 2023. In Proceedings of the 6th Annual ACM Lifelog Search Challenge, pages 53–58. Association for Computing Machinery, New York, NY, USA, 2023.
- [35] L. N. Signal, J. Stanley, M. Smith, M. Barr, T. J. Chambers, J. Zhou, A. Duane, C. Gurrin, A. F. Smeaton, C. McKerchar, et al. Children’s everyday exposure to food marketing: an objective analysis using wearable cameras. International Journal of Behavioral Nutrition and Physical Activity, 14:1–11, 2017.
- [36] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
- [37] F. Spiess, R. Gasser, H. Schuldt, and L. Rossetto. The best of both worlds: Lifelog retrieval with a desktop-virtual reality hybrid system. In Proceedings of the 6th Annual ACM Lifelog Search Challenge, pages 65–68. Association for Computing Machinery, New York, NY, USA, 2023.
- [38] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1–9, 2015.
- [39] W.-C. Tan, J. Dwivedi-Yu, Y. Li, L. Mathias, M. Saeidi, J. N. Yan, and A. Halevy. TimelineQA: A benchmark for question answering over timelines. In A. Rogers, J. Boyd-Graber, and N. Okazaki, editors, Findings of the Association for Computational Linguistics: ACL 2023, pages 77–91, Toronto, Canada, July 2023. Association for Computational Linguistics.
- [40] L.-D. Tran, N. Alam, Y. Graham, L. K. Vo, N. T. Diep, B. Nguyen, L. Zhou, and C. Gurrin. An exploration into the benefits of the CLIP model for lifelog retrieval. In Proceedings of the 19th International Conference on Content-based Multimedia Indexing, pages 15–22, 2022.
- [41] L.-D. Tran, B. Nguyen, L. Zhou, and C. Gurrin. MyEachtra: Event-based interactive lifelog retrieval system for lsc’23. In Proceedings of the 6th Annual ACM Lifelog Search Challenge, pages 24–29. Association for Computing Machinery, New York, NY, USA, 2023.
- [42] L.-D. Tran, M.-D. Nguyen, D.-T. Dang-Nguyen, S. Heller, F. Spiess, J. Lokoč, L. Peška, T.-N. Nguyen, O. S. Khan, A. Duane, et al. Comparing interactive retrieval approaches at the Lifelog Search Challenge 2021. IEEE Access, 11:30982–30995, 2023.
- [43] L.-D. Tran, M.-D. Nguyen, B. Nguyen, H. Lee, L. Zhou, and C. Gurrin. E-Myscéal: Embedding-based interactive lifelog retrieval system for lsc’22. In Proceedings of the 5th Annual on Lifelog Search Challenge, LSC ’22, page 32–37, New York, NY, USA, 2022. Association for Computing Machinery.
- [44] L.-D. Tran, D. Nie, L. Zhou, B. Nguyen, and C. Gurrin. VAISL: Visual-aware identification of semantic locations in lifelog. In International Conference on Multimedia Modeling, pages 659–670. Springer, 2023.
- [45] M.-T. Tran, T.-D. Truong, T. D. Duy, V.-K. Vo-Ho, Q.-A. Luong, and V.-T. Nguyen. Lifelog moment retrieval with visual concept fusion and text-based query expansion. In CLEF (Working Notes), 2018.
- [46] L. Tuovinen and A. F. Smeaton. Privacy-aware sharing and collaborative analysis of personal wellness data: Process model, domain ontology, software system and user trial. Plos ONE, 17(4):e0265997, 2022.

- [47] D. Vrandečić and M. Krötzsch. Wikidata: a free collaborative knowledgebase. Communications of the ACM, 57(10):78–85, 2014.
- [48] G. Wilson, D. Jones, P. Schofield, and D. J. Martin. The use of a wearable camera to explore daily functioning of older adults living with persistent pain: Methodological reflections and recommendations. Journal of Rehabilitation and Assistive Technologies Engineering, 5:2055668318765411, 2018.
- [49] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. Advances in Neural Information Processing Systems, 27, 2014.
- [50] L. Zhou, D.-T. Dang-Nguyen, and C. Gurrin. A baseline search engine for personal life archives. In Proceedings of the 2nd Workshop on Lifelogging Tools and Applications, LTA '17, pages 21–24, New York, NY, USA, 2017. Association for Computing Machinery.
- [51] Q. Zhou, D. Wang, C. Ní Mhurchu, C. Gurrin, J. Zhou, Y. Cheng, and H. Wang. The use of wearable cameras in assessing children’s dietary intake and behaviours in China. Appetite, 139:1–7, 2019.