

DeepDiary: Automatically Captioning Lifelogging Image Streams

Chenyong Fan

David J. Crandall

School of Informatics and Computing
Indiana University
Bloomington, Indiana USA
{fan6,djcran}@indiana.edu

Abstract. Lifelogging cameras capture everyday life from a first-person perspective, but generate so much data that it is hard for users to browse and organize their image collections effectively. In this paper, we propose to use automatic image captioning algorithms to generate textual representations of these collections. We develop and explore novel techniques based on deep learning to generate captions for both individual images and image streams, using temporal consistency constraints to create summaries that are both more compact and less noisy. We evaluate our techniques with quantitative and qualitative results, and apply captioning to an image retrieval application for finding potentially private images. Our results suggest that our automatic captioning algorithms, while imperfect, may work well enough to help users manage lifelogging photo collections.¹

Keywords: Lifelogging; first-person; image captioning; computer vision.

1 Introduction

Wearable cameras that capture first-person views of people’s daily lives have recently become affordable, lightweight, and practical, after many years of being explored only in the research community [1, 12, 22]. These new devices come in various types and styles, from the GoPro, which is marketed for recording high-quality video of sports and other adventures, to Google Glass, which is a heads-up display interface for smartphones but includes a camera, to Narrative Clip and Autographer, which capture “lifelogs” by automatically taking photos throughout one’s day (e.g., every 30 seconds). No matter the purpose, however, all of these devices can record huge amounts of imagery, which makes it difficult for users to organize and browse their image data.

In this paper, we attempt to produce automatic textual *narrations* or captions of a visual lifelog. We believe that describing lifelogs with sentences is most natural for the average user, and allows for interesting applications like generating

¹ An expanded version of this paper is available at <https://arxiv.org/abs/1608.03819>.

automatic textual diaries of the “story” of someone’s day based on their lifelogging photos. We take advantage of recent breakthroughs in image captioning using deep learning that have shown impressive results for consumer-style images from social media [14, 16], and evaluate their performance on the novel domain of first-person images (which are significantly more challenging due to substantial noise, blurring, poor composition, etc.). We also propose a new strategy to try to encourage diversity in the sentences, which we found to be particularly useful in describing lifelogging images from different perspectives.

Of course, lifelogging photo streams are highly redundant since wearable cameras indiscriminately capture thousands of photos per day. Instead of simply captioning individual images, we also consider the novel problem of jointly captioning lifelogging streams, i.e. generating captions for temporally-contiguous groups of photos corresponding to coherent activities or scene types. Not only does this produce a more compact and potentially useful organization of a user’s photo collection, but it also could create an automatically-generated textual “diary” of a user’s day based only on their photos. The sentences themselves are also useful to aid in image retrieval by keyword search, which we illustrate for the specific application of searching for potentially private images (e.g. containing keywords like “bathroom”). We formulate this joint captioning problem in a Markov Random Field model and show how to solve it efficiently.

To our knowledge, we are the first to propose image captioning as an important task for lifelogging photos, as well as the first to apply and evaluate automatic image captioning models in this domain. To summarize our contributions, we learn and apply deep image captioning models to lifelogging photos, including proposing a novel method for generating photo descriptions with diverse structures and perspectives; propose a novel technique for inferring captions for streams of photos taken over time in order to find and summarize coherent activities and other groups of photos; create an online framework for collecting and annotating lifelogging images, and use it to collect a realistic lifelogging dataset consisting of thousands of photos and thousands of reference sentences and evaluate these techniques on our data, both quantitatively and qualitatively, under different simulated use cases.

2 Related Work

While wearable cameras have been studied for over a decade in the research community [1, 12, 22], only recently have they become practical enough for consumers to use on a daily basis. In the computer vision field, recent work has begun to study this new style of imagery, which is significantly different from photos taken by traditional point-and-shoot cameras. Specific research topics have included recognizing objects [8, 18], scenes [9], and activities [4, 7, 25, 26]. Some computer vision work has specifically tried to address privacy concerns, by recognizing photos taken in potentially sensitive places like bathrooms [29], or containing sensitive objects like computer monitors [18]. However, these

techniques typically require that classifiers be explicitly trained for each object, scene type, or activity of interest, which limits their scalability.

Instead of classifying lifelogging images into pre-defined and discrete categories, we propose to annotate them with automatically-generated, free-form image captions, inspired by recent progress in deep learning. Convolutional Neural Networks (CNNs) have recently emerged as powerful models for object recognition in computer vision [6, 10, 19, 28], while Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTMs) have been developed for learning models of sequential data, like natural language sentences [5, 11]. The combination of CNNs for recognizing image content and RNNs for modeling language have recently been shown to generate surprisingly rich image descriptions [14, 23, 32], essentially “translating” from image features to English sentences [15].

Some closely related work has been done to generate textual descriptions from videos. Venugopalan *et al.* [31] use image captioning model to generate video descriptions from a sequence of video frames. Like previous image captioning papers, their method estimates a single sentence for each sequence, while we explicitly generate multiple diverse sentences and evaluate the image-sentence matching quality to improve the captions from noisy, poorly-composed lifelogging images. Zhu *et al.* [33] use neural sentence embedding to model a sentence-sentence similarity function, and use LSTMs to model image-sentence similarity in order to align subtitles of movies with sentences from the original books. Their main purpose is to find corresponding movie clips and book paragraphs based on visual and semantic patterns, whereas ours is to infer novel sentences from new lifelogging image streams.

3 Lifelogging Data Collection

To train and test our techniques, two of the authors wore Narrative Clip lifelogging cameras over a period of about five months (June-Aug 2015 and Jan-Feb 2016), to create a repository of 7,716 lifelogging photos. To facilitate collecting lifelogging photos and annotations, we built a website which allowed users to upload and label photos in a unified framework, using the Narrative Clip API.²

We collected textual annotations for training and testing the system in two different ways. First, the two authors and three of their friends and family members used the online system to submit sentences for randomly-selected images, producing 2,683 sentences for 696 images. Annotators were asked to produce at least two sentences per image: one that described the photo from a first-person perspective (e.g., “I am eating cereal at the kitchen table.”) and one from a third-person perspective (e.g. “A bowl of cereal sits on a kitchen table.”). We requested sentences from each of these perspectives because we have observed that some scenes are more naturally described by one perspective or the other. Annotators were welcome to enter multiple sentences, and each image was viewed by an average of 1.45 labelers.

² <https://open-staging.getnarrative.com/api-docs>

Second, to generate more diversity in annotators and annotations, we published 293 images³ on Amazon’s Mechanical Turk (AMT), showing each photo to at least three annotators and, as before, asking each annotator to give at least one first-person and one third-person sentence. This produced a set of 1,813 sentences, or an average of 6.2 sentences per image. A total of 121 distinct Mechanical Turk users contributed sentences.

Finally, we also downloaded COCO [21], a popular publicly-available dataset of 80,000 photos and 400,000 sentences. These images are from Internet and social media sources, and thus are significantly different than the lifelogging context we consider here, but we hypothesized that this may be useful additional training data to augment our smaller lifelogging dataset.

4 Automatic Lifelogging Image Captioning

We now present our technique for using deep learning to automatically annotate lifelogging images with captions. We first give a brief review of deep image captioning models, and then show how to take advantage of *streams* of lifelogging images by estimating captions jointly across time, which not only helps reduce noise in captions by enforcing temporal consistency, but also helps summarize large photo collections with smaller subsets of sentences.

4.1 Background: Deep networks for image captioning

Automatic image captioning is a difficult task because it requires not only identifying important objects and actions, but also describing them in natural language. However, recent work in deep learning has demonstrated impressive results in generating image and video descriptions [14,31,33]. The basic high-level idea is to learn a common feature space that is shared by both images and words. Then, given a new image, we generate sentences that are “nearby” in the same feature space. The encoder (mapping from image to feature space) is typically a Convolutional Neural Network (CNN), which abstracts images into a vector of local and global appearance features. The decoder (mapping from feature space to words) produces a word vector using a Recurrent Neural Network (RNN), which abstracts out the semantic and syntactic meaning.

In the prediction stage, a forward pass of LSTM generates a full sentence terminated by a stop word for each input image. Similar image captioning models have been discussed in detail in recent papers [14,31,32]. In section 4.2, we discuss in detail how to generate diverse captions for a single image.

4.2 Photo Grouping and Activity Summarization

The techniques in the last section automatically estimate captions for individual images. However, lifelogging users do not typically capture individual images in

³ We randomly chose 300, but removed 7 that we were not comfortable sharing with the public (e.g. photos of strangers whose permission we were not able to obtain).

isolation, but instead collect long streams of photos taken at regular intervals over time (e.g., every 30 seconds for Narrative Clip). This means that evidence from multiple images can be combined together to produce better captions than is possible from observing any single image, in effect “smoothing out” noise in any particular image by examining the photos taken nearby in time. These sentences could provide more concise summarizations, helping people find, remember, and organize photos according to broad events instead of individual moments.

Suppose we wish to estimate captions for a stream of images $I = (I_1, I_2, \dots, I_K)$, which are sorted in order of increasing timestamps. We first generate multiple diverse captions for each individual image, using a technique we describe in the next subsection. We combine all of these sentences together across images into a large set of candidates C (with $|C| = d|I|$, where d is the number of diverse sentences generated per image; we use $d = 15$). We wish to estimate a sequence of sentences such that each sentence describes its corresponding image well, but also such that the sentences are relatively consistent across time. In other words, we want to estimate a sequence of sentences $S^* = (S_1^*, S_2^*, \dots, S_K^*)$ so as to minimize an energy function,

$$S^* = \underset{S=(S_1, \dots, S_K)}{\operatorname{argmin}} \sum_{i=1}^K \operatorname{Score}(S_i, I_i) + \beta \sum_{j=1}^{K-1} \mathbb{1}(S_j, S_{j+1}), \quad (1)$$

where each $S_i \in C$, $\operatorname{Score}(S_i, I_i)$ is a unary cost function measuring the quality of a given sentence S_i in describing a single image I_i , $\mathbb{1}(S_a, S_b)$ is a pairwise cost function that is 0 if S_a and S_b are the same and 1 otherwise, and β is a constant. Intuitively, β controls the degree of temporal smoothing of the model: when $\beta = 0$, for example, the model simply chooses sentences for each image independently without considering neighboring images in the stream, whereas when β is very large, the model will try to find a single sentence to describe all of the images in the stream.

Equation (1) is a chain-structured Markov Random Field (MRF) model [17], which means that the optimal sequence of sentences S^* can be found efficiently using the Viterbi algorithm. All that remains is to define two key components of the model: (1) a technique for generating multiple, diverse candidate sentences for each image, in order to obtain the candidate sentence set C , and (2) the Score function, which requires a technique for measuring how well a given sentence describes a given image. We now describe these two ingredients in turn.

Generating Diverse Captions. Our joint captioning model above requires a large set of candidate sentences. Many possible sentences can correctly describe any given image, and thus it is desirable for the automatic image captioning algorithm to generate multiple sentences that describe the image in multiple ways. This is especially true for lifelogging images that are often noisy, poorly composed, and ambiguous, and can be interpreted in different ways. Vinyals *et al.* [32] use beam search to generate multiple sentences, by having the LSTM model keep b candidate sentences at each step of sentence generation (where



Fig. 1: Sample captions generated by models pre-trained with COCO and fine-tuned with lifelogging dataset. Three different colors show the top three predictions produced in three beam searches by applying the Diverse M-Solutions technique. Within each beam search, sentences tend to have similar structures and describe from similar perspective; between consecutive beam searches, structures and perspectives tend to be different.

b is called the beam size). However, we found that this existing technique did not work well for lifelogging sentences, because it produced very homogeneous sentences, even with a high beam size.

To encourage greater diversity, we apply the Diverse M-best solutions technique of *Batra et al.* [3], which was originally proposed to find multiple high-likelihood solutions in graphical model inference problems. We adapt this technique to LSTMs by performing multiple rounds of beam search. In the first round, we obtain a set of predicted words for each position in the sentence. In the second round, we add a bias term that reduces the network activation values of words found in the first beam search by a constant value. Intuitively, this decreases the probability that a word found during the previous beam search being selected again at the same word position in the sentence. Depending on the degree of diversity needed, additional rounds of beam search can be conducted, each time penalizing words that have occurred in any previous round.

Figure 1 presents sample automatically-generated results by using three rounds of beam search and a beam size of 3 for illustration purposes. We see that the technique successfully injects diversity into the set of estimated captions. Many of the captions are quite accurate, including “A man is sitting at a table” and “I am having dinner with my friends,” while others are not correct (e.g. “A man is looking at a man in a red shirt”), and others are nonsensical (“There

is a man sitting across the table with a man”). Nevertheless, the captioning results are overall remarkably accurate for an automatic image captioning system, reflecting the power of deep captioning techniques to successfully model both image content and sentence generation.

Image-sentence quality alignment. The joint captioning model in Equation (1) also requires a function $\text{Score}(S_i, I_i)$, which is a measure of how well an arbitrary sentence S_i describes a given image I_i . The difficulty here is that the LSTM model described above tells us how to generate sentences for an image, but not how to measure their similarity to a given image. Doing this requires us to explicitly align certain words of the sentence to certain regions of an image – i.e. determining which “part” of an image generated each word. Karpathy *et al.* [16] propose matching each region with the word with maximum inner product (interpreted as a similarity measure) across all words in terms of learnable region vectors and word vectors, and to sum all similarity measures over all regions as the total score. We implement their method and train this image-sentence alignment model on our lifelogging dataset. To generate the matching score $\text{Score}(S_i, I_i)$ for Equation (1), we extract region vectors from image I_i , retrieve trained word vectors for words in sentence S_i , and sum similarity measures of regions with best-aligned words.

Image Grouping Result. Finally, once captions have been jointly inferred for each image in a photostream, we can group together contiguous substreams of images that share the same sentence. Figure 2 shows examples of activity summarization. In general, the jointly-inferred captions are reasonable descriptions of the images, and much less noisy than those produced from individual images in Figure 1, showing the advantage of incorporating temporal reasoning into the captioning process. For example, the first row of images shows that the model labeled several images as “I am talking with a friend while eating a meal in a restaurant,” even though the friend is only visible in one of the frames, showing how the model has propagated context across time. Of course, there are still mistakes ranging from the minor error that there is no broccoli on the plate in the second row to the more major error that the last row shows a piano and not someone typing on a computer. The grammar of the sentences is generally good considering that the model has no explicit knowledge of English besides what it has learned from training data, although usage errors are common (e.g., “I am shopping kitchen devices in a store”).

5 Experimental evaluation

We first use automatic metrics that compare to ground truth reference sentences with quantitative scores. To give a better idea of the actual practical utility of technique, we also evaluate in two other ways: using a panel of human judges to rate the quality of captioning results, and testing the system in a specific application of keyword-based image retrieval using the generated captions.

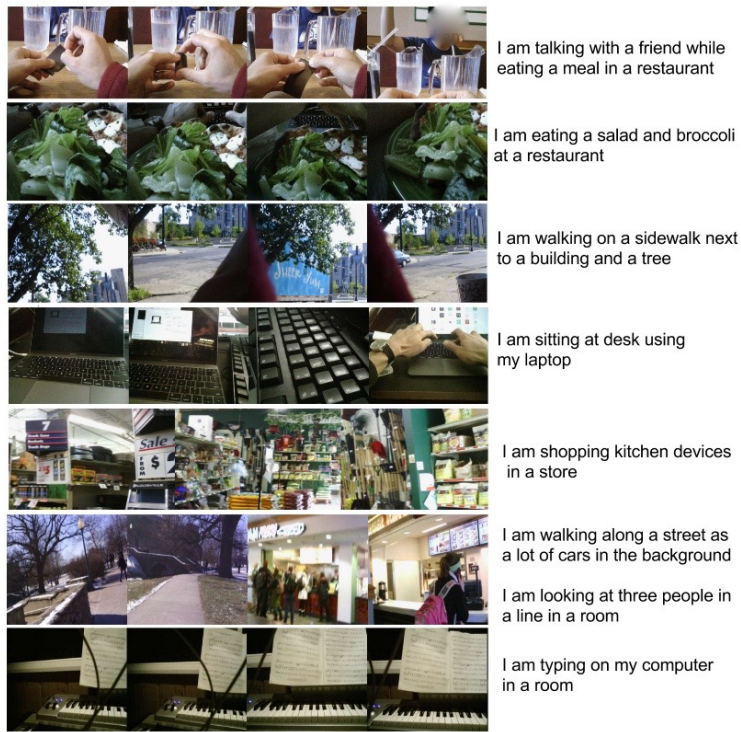


Fig. 2: Randomly-chosen samples of activity summarization on our dataset.

5.1 Quantitative captioning evaluation

Automated metrics such as BLEU [24], CIDEr [30], Meteor [2] and Rouge-L [20] have been proposed to score sentence similarity compared to reference sentences provided by humans, and each has different advantages and disadvantages. We present results using all of these metrics (using the MS COCO Detection Challenge implementation⁴), and also summarize the seven scores with their mean.

Implementation. A significant challenge with deep learning-based methods is that they typically require huge amounts of training data, both in terms of images and sentences. Unfortunately, collecting this quantity of lifelogging images and annotations is very difficult. To try to overcome this problem, we augmented our lifelogging training set with COCO data using three different strategies: **Lifelog only** training used only our lifelogging dataset, consisting of 736 lifelogging photos with 4,300 human-labeled sentences; **COCO only** training used only COCO dataset; and **COCO then Lifelog** started with the **COCO only** model,

⁴ <https://github.com/tylin/coco-caption>

Table 1: Bleu1-4, CIDEr, Meteor and Rouge Scores for Diverse 3-Best Beams of Captions on Test Set.

<i>Datasets</i>		<i>Metric</i>							
Training	Testing	Bleu-1	Bleu-2	Bleu-3	Bleu-4	CIDEr	METEOR	ROUGE	Mean
Lifelog		0.669	0.472	0.324	0.218	0.257	0.209	0.462	0.373
COCO	Lifelog 100	0.561	0.354	0.206	0.118	0.143	0.149	0.374	0.272
COCO+Lifelog		0.666	0.469	0.319	0.210	0.253	0.207	0.459	0.369
Lifelog@Usr1	Lifelog@Usr2	0.588	0.410	0.279	0.189	0.228	0.195	0.431	0.331
Lifelog@2015	Lifelog@2016	0.557	0.379	0.249	0.160	0.325	0.202	0.425	0.328

and then used it as initialization when re-training the model on the lifelogging dataset (i.e., “fine-tuning” [19]).

For extracting image features, we use the VGGNet [27] CNN model. The word vectors are learned from scratch. Our image captioning model stacks two LSTM layers, and each layer structure closely follows the one described in [32]. To boost training speed, we re-implemented LSTM model in C++ using the Caffe [13] deep learning package. It takes about 2.5 hours for COCO pre-training, and about 1 hour for fine-tuning on Lifelog dataset with 10,000 iterations for both.

At test time, the number of beam searches conducted during caption inference controls the degree of diversity in the output; here we use three to match the three styles of captions we expect (COCO, first-person, and third-person perspectives). Samples of predicted sentences are shown in Figure 1. This suggests that different genres of training sentences contribute to tune hidden states of LSTM and thus enable it to produce diverse structures of sentences in testing stage.

Results. Table 1 presents quantitative results of each of these training strategies, all tested on the same set of 100 randomly-selected photos having 1,000 ground truth reference sentences, using each of the seven automatic scoring metrics mentioned above. We find that the **Lifelog only** strategy achieves much higher overall accuracy than **COCO only**, with a mean score of 0.373 vs. 0.272. This suggests that even though COCO is a much larger dataset, images from social media are different enough from lifelogging images that the **COCO only** model does not generalize well to our application. Moreover, this may also reflect an artifact of the automated evaluation, because **Lifelog only** benefits from seeing sentences with similar vocabulary and in a similar style as in the reference sentences, since the same small group of humans labeled both the training and test datasets. More surprisingly, we find that **Lifelog only** also slightly outperforms **COCO then lifelog** (0.373 vs 0.369). The model produced by the latter training dataset has a larger vocabulary and produces richer styles of sentences than Lifelog only, which hurts its quantitative score. Qualitatively, however, it often produces more diverse and descriptive sentences because of its larger vocabulary and ability to generate sentences in first-person, third-person, and COCO styles. Samples of generated diverse captions are shown in Figure 1.

We conducted experiments with two additional strategies in order to simulate more realistic scenarios. The first scenario reflects when a consumer first starts using our automatic captioning system on their images without having supplied any training data of their own. We simulate this by training image captioning model on one user’s photos and testing on another. Training set has 805 photos and 3,716 reference sentences; testing set has 40 photos and 565 reference sentences. The mean quantitative accuracy declines from our earlier experiments when training and testing on images sampled from the same set, as shown in Table 1, although the decline is not very dramatic (from 0.373 to 0.331), and still much better than training on COCO (0.272). This result suggests that the captioning model has learned general properties of lifelogging images, instead of overfitting to one particular user (e.g., simply “memorizing” the appearance of the places and activities they frequently visit and do).

The other situation is when an existing model trained on historical lifelogging data is used to caption new photos. We simulate this by taking all lifelogging photos in 2015 as training data and photos in 2016 as testing data. Training set has 673 photos and 3,610 sentences; testing set has 30 photos and 172 sentences. As shown in Table 1, this scenario very slightly decreased performance compared to training on data from a different user (0.328 vs 0.331), although the difference is likely not statistically significant.

5.2 Image captioning evaluation with human judges

We conducted a small study using human judges to rate the quality of our automatically-generated captions. In particular, we randomly selected 21 images from the Lifelog 100 test dataset (used in Table 1) and generated captions using our model trained on the COCO then Lifelog scenario. For each image, we generated 15 captions (with 3 rounds of beam search, each with beam size 5), and then kept the top-scoring caption according to our model and four randomly-sampled from the remaining 14, to produce a diverse set of five automatically-generated captions per image. We also randomly sampled five of the human-generated reference sentences for each image.

For each of the ten captions (five automatic plus five human), we showed the image (after reviewing it for potentially private content and obtaining permission of the photo-taker) and caption to a user on Amazon Mechanical Turk, without telling them how the caption had been produced. We asked them to rate, on a five-point Likert scale, how strongly they agreed with two statements: (1) “The sentence or phrase makes sense and is grammatically correct (ignoring minor problems like capitalization and punctuation,” and (2) “The sentence or phrase accurately describes either what the camera wearer was doing or what he or she was looking at when the photo was taken.” The task involved 630 individual HITs from 37 users.

Table 2 summarizes the results, comparing the average ratings over the 5 human reference sentences, the average over all 5 diverse automatically-generated captions (Auto-5 column), and the single highest-likelihood caption as estimated by our complete model (Auto-top). About 92% of the human reference sentences

Table 2: Summary of grammatical correctness and accuracy of lifelogging image captions, on a rating scale from 1 (Strongly Disagree) to 5 (Strongly Agree), averaged over 3 judges. *Human* column is averaged over 5 human-generated reference sentences, *Auto-5* is averaged over 5 diverse computer-generated sentences, and *Auto-top* is single highest-likelihood computer-generated sentence predicted by our model.

Rating	Grammar			Accuracy		
	Human	Auto-5	Auto-top	Human	Auto-5	Auto-top
1	1.9%	7.6%	11.9%	2.9%	22.4%	21.4%
2	3.8%	10.0%	7.1%	3.8%	15.2%	7.1%
3	0.5%	5.7%	0.0%	4.8%	8.1%	7.1%
4	19.0%	17.6%	4.8%	22.4%	17.6%	19.0%
5	73.3%	59.0%	76.2%	65.2%	36.7%	45.2%
Mean	4.60	4.10	4.26	4.45	3.31	3.60

were judged as grammatically correct (i.e., somewhat or strongly agreeing with statement (1)), compared to about 77% for the automatically-generated diverse captions and 81% for the single best sentence selected by our model. Humans also described images more accurately than the diverse captions (88% vs 54%), although the fact that 64% of our single best estimated captions were accurate indicates that our model is often able to identify which one is best among the diverse candidates. Overall, our top automatic caption was judged to be *both* grammatically correct and accurate 59.5% of the time, compared to 84.8% of the time for human reference sentences.

We view these results to be very promising, as they suggest that automatic captioning can generate reasonable sentences for over half of lifelogging images, at least in some applications. For example, for 19 (90%) of the 21 images in the test set, at least one of five diverse captions was unanimously judged to be both grammatically correct and accurate by all 3 judges. This may be useful in some retrieval applications where recall is important, for example, where having noise in some captions may be tolerable as long as at least one of them is correct. We consider one such application in the next section.

5.3 Keyword-based image retrieval

Image captioning allows us to directly implement keyword-based image retrieval by searching on the generated captions. We consider a particular application of this image search feature here that permits a quantitative evaluation. As mentioned above, wearable cameras can collect a large number of images containing private information. Automatic image captioning could allow users to find potentially private images easily, and then take appropriate action (like deleting or encrypting the photos). We consider two specific types of potentially embarrassing content here: photos taken in potentially private locations like bathrooms and locker rooms, and photos containing personal computer or smartphone displays which may contain private information such as credit card numbers or e-mail contents.

Table 3: Confusion matrices for two approaches on two tasks for detecting sensitive images. *Left*: Results on 3-way problem of classifying into not sensitive, sensitive place (bathroom), or digital display categories. *Right*: Results on 2-way problem of classifying into sensitive or not (regardless of sensitivity type). Actual classes are in rows and predicted classes are in columns.

	3-way classification						2-way classification				
	CNN-based			Caption-based			CNN-based		Caption-based		
	NotSen	Place	Display	NotSen	Place	Display	NotSen	Sen	NotSen	Sen	
NotSen	0.730	0.130	0.140	0.686	0.117	0.197	NotSen	0.730	0.270	0.686	0.314
Place	0.189	0.811	0	0.151	0.792	0.057	Sen	0.317	0.683	0.161	0.839
Display	0.300	0.043	0.657	0.143	0.008	0.849					

We chose these two types of concerns specifically because they have been considered by others in prior work: Korayem *et al.* [18] present a system for detecting monitors in lifelogging images using deep learning with CNNs, while Templeman *et al.* [29] classify images according to the room in which they were taken. Both of these papers present strongly supervised based techniques, which were given thousands of training images manually labeled with ground truth for each particular task. In contrast, identifying private imagery based on keyword search on automatically-generated captions could avoid the need to create a training set and train a separate classifier for each type of sensitive image.

We evaluated captioning-based sensitive image retrieval against standard state-of-the-art strongly-supervised image classification using CNNs [19] (although we cannot compare directly to the results presented in [18] or [29] because we use different datasets). We trained the strongly-supervised model by first generating a training set consisting of photos having monitors and not having monitors, and photos taken in bathrooms and locker rooms or elsewhere, by using the ground truth categories given in the COCO and Flickr8k datasets. This yielded 34,736 non-sensitive images, 6,135 images taken in sensitive places, and 4,379 images with displays. We used pre-trained AlexNet model (1000-way classifier on ImageNet data) and fine-tuned on our dataset by replacing the final fully connected layer with a 3-way classifier to correspond with our three-class problem.

We also ran the technique proposed here, where we first generate automatic image captions, and then search through the top five captions for each image for a set of pre-defined keywords (specifically “toilet,” “bathroom,” “locker,” “lavatory,” and “washroom” for sensitive place detection, and “computer,” “laptop,” “iphone,” “smartphone,” and “screen” for display detection). If any of these keywords is detected in any of the five captions, the image is classified as sensitive, and otherwise it is estimated to be not sensitive.

Table 3 presents the confusion matrix for each method, using a set of 600 manually-annotated images from our lifelogging dataset as test data (with 300 non-sensitive images, 53 images in sensitive places, and 252 with digital displays). We see the supervised classifier has better prediction performance on finding sensitive places (0.811) than keyword based classifiers (0.792), while the caption-based technique classifier outperforms on predicting second type of sensitive




Fig. 3: Precision-recall curves for retrieving sensitive images using CNNs (left) and generated captions (right).

images (0.849 vs 0.657). In a real application, determining the type of private image is likely less important than simply deciding if it is private. The bottom table in Table 3 reflects this scenario, showing a confusion matrix which combines the two sensitive types and focuses on whether photos are sensitive or not.

From another point of view, sensitive photo detection is a retrieval problem. Figure 3 shows precision-recall curves for CNN and caption-based classifiers, respectively. They show the trade-off between selecting accurate sensitive photos (high precision) and obtaining a majority of all sensitive photos (high recall). For example, by using CNN classifier, we can obtain 80% type 1 (sensitive place) photos with accuracy around 58% (Figure 3(left) green curve); by using the caption-based classifier, we can obtain 80% of type 2 (digital display) sensitive photos with precision around 78% (Figure 3(right) blue curve).

The two approaches may also be complementary, since they use different forms of evidence in making classification decisions, and users in a real application could choose their own trade-off on how aggressively to filter lifelogging images.

6 Conclusion

In this paper, we have proposed the concept of using automatically-generated captions to help organize and annotate lifelogging image collections. We have proposed a deep learning-based captioning model that jointly labels photo streams in order to take advantage of temporal consistency between photos. Our evaluation suggests that modern automated captioning techniques could work well enough to be used in practical lifelogging photo applications. We hope our research will motivate further efforts of using lifelogging photos and descriptions together to help human memory recall the activities and scenarios.

Acknowledgements. This work was supported in part by the National Science Foundation (IIS-1253549 and CNS-1408730) and Google, and used compute facilities provided by NVidia, the Lilly Endowment through support of the IU PTI, and the Indiana METACyt Initiative. We thank Zhenhua Chen, Sally Crandall, and Xuan Dong for helping to label our lifelogging photos.

References

1. Azuma, R., Baillot, Y., Behringer, R., Feiner, S., Julier, S., MacIntyre, B.: Recent advances in augmented reality. *IEEE Computer Graphics and Applications* 21(6), 34–47 (2001)
2. Banerjee, S., Lavie, A.: METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. *ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization* pp. 65–72 (2005)
3. Batra, D., Yadollahpour, P., Guzman-Rivera, A., Shakhnarovich, G.: Diverse m-best solutions in markov random fields. In: *European Conf. on Computer Vision*, pp. 1–16. Springer (2012)
4. Castro, D., Hickson, S., Bettadapura, V., Thomaz, E., Abowd, G., Christensen, H., Essa, I.: Predicting daily activities from egocentric images using deep learning. In: *Intl. Symposium on Wearable Computers* (2015)
5. Elman, J.L.: Finding structure in time. *Cognitive science* 14(2), 179–211 (1990)
6. Erhan, D., Szegedy, C., Toshev, A., Anguelov, D.: Scalable object detection using deep neural networks. In: *IEEE Conf. on Computer Vision and Pattern Recognition*. pp. 2155–2162 (2014)
7. Fathi, A., Li, Y., Rehg, J.M.: Learning to recognize daily actions using gaze. In: *European Conf. on Computer Vision*, pp. 314–327. Springer (2012)
8. Fathi, A., Ren, X., Rehg, J.M.: Learning to recognize objects in egocentric activities. In: *IEEE Conf. on Computer Vision and Pattern Recognition*. pp. 3281–3288 (2011)
9. Furnari, A., Farinella, G., Battiano, S.: Recognizing personal contexts from egocentric images. In: *ICCV Workshops* (2015)
10. Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: *IEEE Conf. on Computer Vision and Pattern Recognition*. pp. 580–587 (2014)
11. Graves, A.: Generating sequences with recurrent neural networks. [arXiv:1308.0850](https://arxiv.org/abs/1308.0850) (2013)
12. Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., Wood, K.: Sensecam: A retrospective memory aid. In: *ACM Conf. on Ubiquitous Computing*. pp. 177–193 (2006)
13. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: Convolutional architecture for fast feature embedding. [arXiv:1408.5093](https://arxiv.org/abs/1408.5093) (2014)
14. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. [arXiv:1412.2306](https://arxiv.org/abs/1412.2306) (2014)
15. Karpathy, A., Johnson, J., Fei-Fei, L.: Visualizing and understanding recurrent networks. [arXiv:1506.02078](https://arxiv.org/abs/1506.02078) (2015)
16. Karpathy, A., Joulin, A., Fei-Fei, L.: Deep fragment embeddings for bidirectional image sentence mapping. In: *Advances in neural information processing systems*. pp. 1889–1897 (2014)
17. Koller, D., Friedman, N.: *Probabilistic Graphical Models: Principles and Techniques*. MIT Press (2009)
18. Korayem, M., Templeman, R., Chen, D., Crandall, D., Kapadia, A.: Enhancing lifelogging privacy by detecting screens. In: *ACM CHI Conf. on Human Factors in Computing Systems* (2016)
19. Krizhevsky, A., Sutskever, I., Hinton, G.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*. pp. 1097–1105 (2012)

20. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: Workshop On Text Summarization Branches Out (2004)
21. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: Common objects in context. In: European Conf. on Computer Vision, pp. 740–755. Springer (2014)
22. Mann, S., Nolan, J., Wellman, B.: Sousveillance: Inventing and using wearable computing devices for data collection in surveillance environments. *Surveillance & Society* 1(3), 331–355 (2002)
23. Mao, J., Xu, W., Yang, Y., Wang, J., Yuille, A.L.: Explain images with multimodal recurrent neural networks. arXiv:1410.1090 (2014)
24. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: BLEU: a method for automatic evaluation of machine translation. In: Annual Meeting of the Association for Computational Linguistics. pp. 311–318 (2002)
25. Ryoo, M., Matthies, L.: First-person activity recognition: What are they doing to me? In: IEEE Conf. on Computer Vision and Pattern Recognition. pp. 2730–2737 (2013)
26. Ryoo, M., Fuchs, T.J., Xia, L., Aggarwal, J.K., Matthies, L.: Robot-centric activity prediction from first-person videos: What will they do to me. In: ACM/IEEE Intl. Conf. on Human-Robot Interaction. pp. 295–302 (2015)
27. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv:1409.1556 (2014)
28. Szegedy, C., Toshev, A., Erhan, D.: Deep neural networks for object detection. In: Advances in Neural Information Processing Systems. pp. 2553–2561 (2013)
29. Templeman, R., Korayem, M., Crandall, D.J., Kapadia, A.: Placeavoider: Steering first-person cameras away from sensitive spaces. In: Network and Distributed Systems Security Symposium (2014)
30. Vedantam, R., Zitnick, C., Parikh, D.: Cider: Consensus-based image description evaluation. In: IEEE Conf. on Computer Vision and Pattern Recognition. pp. 4566–4575 (2015)
31. Venugopalan, S., Rohrbach, M., Donahue, J., Mooney, R., Darrell, T., Saenko, K.: Sequence to sequence–video to text. arXiv:1505.00487 (2015)
32. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. arXiv:1411.4555 (2014)
33. Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., Fidler, S.: Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. arXiv:1506.06724 (2015)