arXiv:1712.09915v1 [cs.MM] 28 Dec 2017

Field Studies with Multimedia Big Data: Opportunities and Challenges (Extended Version)

Mario Michael Krell, Julia Bernd, Yifan Li, Daniel Ma , Jaeyoung Choi, Michael Ellsworth, Damian Borth, and Gerald Friedland

TR-17-002

December 2017

FIELD STUDIES WITH MULTIMEDIA BIG DATA: OPPORTUNITIES AND CHALLENGES (EXTENDED VERSION)

Mario Michael Krell[1,2,*], *Julia Bernd*[1,*], *Yifan Li*[2], *Daniel Ma*[2], *Jaeyoung Choi*[1,3], *Michael Ellsworth*[1,2], *Damian Borth*[4], *Gerald Friedland*[1,2,5]
[1]International Computer Science Institute, Berkeley, [2]University of California, Berkeley, [3]Delft University of Technology, Delft, [4]German Research Center for Artificial Intelligence, Saarbrücken, [5]Lawrence Livermore National Laboratory, Livermore. *These authors contributed equally to this work.

ABSTRACT

Social multimedia users are increasingly sharing all kinds of data about the world. They do this for their own reasons, not to provide data for field studies—but the trend presents a great opportunity for scientists. The Yahoo Flickr Creative Commons 100 Million (YFCC100M) dataset comprises 99 million images and nearly 800 thousand videos from Flickr, all shared under Creative Commons licenses. To enable scientists to leverage these media records for field studies, we propose a new framework that extracts targeted subcorpora from the YFCC100M, in a format usable by researchers who are not experts in big data retrieval and processing.

This paper discusses a number of examples from the literature—as well as some entirely new ideas—of natural and social science field studies that could be piloted, supplemented, replicated, or conducted using YFCC100M data. These examples illustrate the need for a general new open-source framework for Multimedia Big Data Field Studies. There is currently a gap between the separate aspects of what multimedia researchers have shown to be possible with consumer-produced big data and the follow-through of creating a *comprehensive* field study framework that supports scientists across other disciplines.

To bridge this gap, we must meet several challenges. For example, the framework must handle unlabeled and noisily labeled data to produce a filtered dataset for a scientist—who naturally wants it to be both as large and as clean as possible. This requires an iterative approach that provides access to statistical summaries and refines the search by constructing new classifiers. The first phase of our framework is available as Multimedia Commons Search (`http://search.mmcommons.org`, MMCS), an intuitive interface that enables complex search queries at a large scale. After outlining our proposal for the general framework and discussing the potential example studies, this paper describes and evaluates a practical application to the study of origami.

1. INTRODUCTION

The basis of science is quite often data. Consequently, data science and machine learning are very hot topics, with new applications and insights coming out every hour. Unfortunately, it can take a lot of time to record or gather that data, and to add the necessary annotations and metadata for machine learning.

For scientific field studies generally, a large proportion of researchers' time is often spent in administrative tasks associated with data-gathering. For example, they must seek funding, get approval from institutional ethics committees or other relevant committees, and, if children are involved, obtain parental consent. Additionally, the data recording itself might take a lot of time, if many samples are needed to confirm or invalidate a hypothesis. In particular, automatic analysis of data using machine learning requires quite large datasets. Last but not least, the sampling variety a scientist can achieve with a given dataset, in terms of geographical and cultural diversity, is generally constrained by the time and money required for travel.

The Yahoo Flickr Creative Commons 100 Million (YFCC100M) is the largest publicly available multimedia dataset [127]. This user-generated content (UGC) corpus comprises 99.2 million images and 800,000 videos from Flickr, shared under Creative Commons copyrighticenses. New extensions and subsets are frequently added, as part of the Multimedia Commons initiative [9]. An important advantage of the YFCC100M for scientific studies is that access is open and simple. Because the media are Creative Commons, they can be used for almost any type of research—and the standardized licenses make any restrictions clear (for example, on commercial applications). In Figure 1 some example images and licenses are shown.

We are proposing a framework to extend the existing YFCC100M ecosystem to enable scientists with no expertise in big data retrieval and processing to filter out data that is relevant for their research question, and to process that data to begin answering it. Depending on the topic, the framework might be used to collect and analyze the central dataset for a new study, or it might be used to pilot or prepare for on-

Fig. 1. YFCC100M images extracted by using the pre-existing data browser (http://YFCC100M.appspot.com/about) [73] to search for 'throw', 'dog baby', and 'parrot' (from left to right: [https://www.flickr.com/photos/29233640@N07/3318234428, license: CC BY 2.0], [https://www.flickr.com/photos/7300793@N06/5725224597, license: CC BY-NC 2.0], [https://www.flickr.com/photos/69221340@N04/6912680630, license: CC BY 2.0]).

the-ground data collection. In addition, the framework can provide a way to replicate or extend existing studies, where the original recordings are not publicly available due to administrative restrictions or ethics rules.

The contributions of this paper are in introducing the multimedia big data studies (MMBDS) concept and our search engine for the YFCC100M, discussing the benefits and limitations of MMBDS, providing many elaborated examples of potential MMBDS from different disciplines, and analyzing some of the requirements for implementing those studies in a comprehensive general framework (including existing tools that could be leveraged). For the requirements analysis, we apply our concept to the concrete and entirely novel example of origami studies.

Overall, we hope the new possibilities offered by this framework will inspire scientists to come up with new research ideas that leverage the Multimedia Commons, and to contribute new tools and data resources to the framework. We thus can bridge the gap between the *potential* applications of multimedia research and its *actual* application in the sciences and humanities.

In Section 2, we describe the YFCC100M dataset ecosystem in more detail, with reference to its potential for MMBDS. In Section 3, we describe the basic structure of the framework. Next, we give examples of studies that could be extended, new research that could be done in future, and new applications that could be pursued using the YFCC100M or similar datasets in an MMBDS framework (Sections 4 and 5). We describe a practical case study on origami in Section 6. Section 7 provides an outlook for the future.

2. THE YFCC100M DATASET

This section describes the YFCC100M ecosystem, including the dataset, associated extensions, and tools, focusing on the advantages and limitations for MMBDS.

While we eventually intend to incorporate additional data sources into the MMBDS framework, we began with the YFCC100M because of its size, its availability, and other suitable characteristics.

2.1. Purpose-Built vs. General Datasets

If a researcher can get all the data required for a study from an existing purpose-built dataset, that is usually the best choice. But quite often, existing datasets might not be diverse enough in terms of location, language, etc. to suit their needs, or they might not have the right data type (text, image, audio, video).

Scraping data from the web presents a different set of difficulties. Existing search engines (whether general, like Google, or service-specific, like Flickr or YouTube) are not built for data-gathering, and are only partially able to handle unlabeled data. Constant updates to the search engines and the content mean that the search that led to the dataset will not necessarily be *reproducible*. Providing the actual data, so other researchers can *replicate* the study, may not be possible due to restrictive or unclear licensing, and maintaining the data in a public repository requires long-term resources.

In contrast, the YFCC100M can be accessed at any time, the data in it remains the same, and a subset can easily be shared as a list of IDs.

2.2. YFCC100M Characteristics

The YFCC100M is comprised of images and videos that were uploaded to Flickr under Creative Commons licenses between

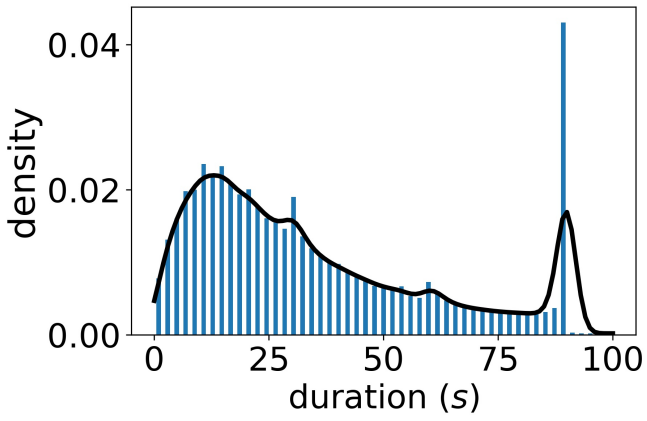

Fig. 2. Duration distribution of YFCC100M videos.

2004 and 2014. This comparatively long time frame—and the enormous amount of data—make it particularly suited to scientific studies.

Classification results show that the YFCC100M is very diverse in terms of subject matter [127]. And again, even if a researcher cannot find a sufficiently large set of high-quality images or videos for a full study on their topic, the YFCC100M can be very helpful for exploratory or preliminary studies. Such pre-studies can help a researcher decide which factors to vary or explore in an in-depth, controlled study. They can also be quite useful in illuminating potential difficulties or confounds. This can help researchers avoid mistakes and delays during targeted data acquisition.

The average length of videos in the YFCC100M is 39 seconds. Some are, of course, much longer—but for studies that require longer continuous recordings, there might not be enough long examples. However, Sections 4 and 5 provide a number of examples where short videos would suffice, such as analyzing human body movements.

In addition to the raw data, the YFCC100M dataset includes metadata such as user-supplied tags and descriptions, locations (GPS coordinates and corresponding place names), recording timestamps, and camera types, for some or all of the media. This metadata can be also used in MMBDS. In particular, location (available for about half the media) could be highly relevant when a study requires data from a specific region or when location is a factor in the study, e.g., for analysis of environmental or cultural differences between regions. Timestamps can be used when changes over time are of interest, such as changes in snow cover.[1] Camera types may influence features of images or videos in ways that are relevant for extraction of information.

On-the-ground field studies are limited in how many specific locations they can target, due to time and resources. In contrast, the wide coverage across locations found in the YFCC100M [28, 73] makes it possible to focus on comparing many different places—or, conversely, to reduce location bias within a study area.

However, metadata—especially when it is user-generated—has its limits. Complete and correct metadata cannot be expected. A robust approach leveraging this wealth of metadata must therefore work around incorrect, ambiguous, or missing annotations. We discuss this in detail in Section 3.1.

[1] GPS coordinates and timestamps are not always accurate, but inaccurate data is usually easy to identify and discard.

2.3. Extended Resources and Tools

An important advantage of the YFCC100M dataset is that new resources are often added by research groups working together in the Multimedia Commons. These include carefully selected annotated subsets [e.g., 10] and preprocessed datasets (or subsets) of commonly used types of features [e.g., 3, 9, 28, 71, 111]. Having precomputed features can significantly speed up processing. The Multimedia Commons also includes a set of automatically generated tags (autotags) for all of the images and the first frame of each video, labeling them for 1500 visual concepts (object classes) with 90% precision [127], and a set of automatically generated multimodal video labels for 609 concepts [71].

The existing subsets can reduce the amount of data that has to be parsed to extract relevant examples for a study. As one example, the YLI-MED video subset [10] provides strong annotations for ten targeted events (or no targeted event), along with attributes like languages spoken and musical scores. Our objective in part is to enable scientists to create new strongly annotated subsets (as described in Section 3) and—ideally—to contribute them in turn to the Multimedia Commons.

The current Multimedia Commons ecosystem includes an easy-to-use image browser, described in Kalkowski et al. 2015 [73]. This browser can use the image metadata to generate subsets according to a user's specifications, provide statistics about that subset or the dataset as a whole, allow users to view the images, and provide URLs to download images for further analysis.

3. THE MMBDS FRAMEWORK

Our MMBDS framework takes a similar approach to Kalkowski et al.'s data browser [73]. However, our new framework is open source, enables more types of searches (e.g., feature-based), and provides more ways to interact with and refine a dataset to achieve the desired result.

This section outlines the MMBDS framework—including specifications based on our conversations with scientists and our case study on origami (see Section 6)—and describes our progress in implementing it.

3.1. The Proposed Dataset-Building Process

Ideally, a scientist will want high-quality data with strong—i.e., consistent and reliable—annotations. But user-supplied tags are generally inconsistent, sometimes inaccurate, and often do not exist at all. In addition, scientists will frequently be looking for characteristics that a user would not typically think to tag because they are unremarkable or backgrounded (like ordinary trees lining a street)—or would not tag in the same way (for example, listing each species of tree). However, obviously, it would be very cumbersome for expert annotators to go through all of the YFCC100M images and videos and label them by hand for a given research project.

We propose an iterative, hybrid approach to take advantage of both content and metadata in this large corpus. A typical search process might start with a user selecting some terms and filters to gather an initial candidate pool. These filters can be of multiple types, including metadata filters and/or (weak) detectors. The use of automatic detectors means that all data is included in the search, whether the user has supplied tags or not. The search engine would enable the user to prioritize the various filters to sort the data [as in 70, 141]. (Alternatively, the user could start with a similarity search, if they already had some relevant examples on hand.)

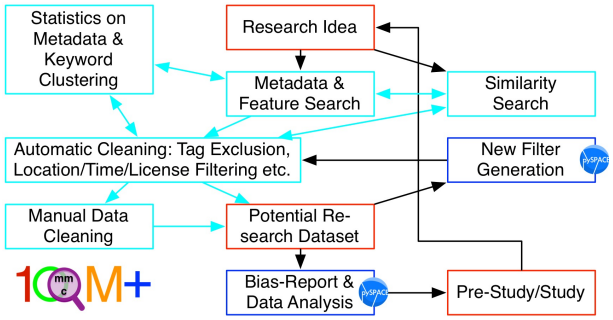

Fig. 3. Schematic representation of the MMBDS framework, including ideas and products (red), dataset extraction (cyan) handled by MMCS and scripts, and data analysis (blue) handled by pySPACE.

If the resulting set of candidates is sufficient and reasonably apropos, the search could end there, with the expert (optionally) manually eliminating the less relevant examples. The data browser may then be used to add any additional annotations the researcher requires. If it is not satisfactory, the candidate pool could be automatically narrowed. Narrowing could be done by adding or removing filters/detector types, changing the filter parameters, adjusting their associated confidence bounds, or selecting some relevant candidates and keeping only the most similar examples among the rest.

Alternatively, if the candidate pool is too small or narrow, the next step could be to expand it. This could be done by adding or removing filters/detector types, adjusting filter parameters or confidence bounds, or using similarity search based on the best candidates found so far. (Where similarity search might be based on content/features or on metadata clustering.) The user could also examine the metadata (including autotags) associated with examples identified via automatic detectors, to inspire new metadata search terms.

Finally, the data processing framework (see Section 3.3) would allow the user to create a new classifier/filter based on the candidate pool, and re-apply it to a new search.

Each of these processes could be iterated until the field expert is satisfied with the size and quality of the example dataset. Although the expert would most likely have to do some manual reviewing and selecting, the automatic filtering would make each step more manageable. In addition, any labels generated in the process could (if the user chose) be fed back into the system, to provide more reliable annotations for future studies. This process is summarized in Figure 3.

3.2. Implementation: Expanding Search Capabilities

There are, of course, a number of existing approaches to search. But so far, none brings together all the features needed for MMBDS. For MMBDS, a scientist should be able to iteratively choose and customize searches on all different types and aspects of multimedia data (text, textual metadata, images, audio, etc.); retrieve data by labels, by content similarity, or by specifying particular characteristics for detection; specify the fuzziness of the parameters; and (if desired) retrieve *all* of the matching examples. In addition to browsing and selecting data, the framework should also allow for annotation and for the creation of new filters within the same interface (see Section 3.3). The testing of new filters, as well as the goal of full customizability, can be much better achieved when the framework is fully open source—unlike the current YFCC100M browser [73].

We are therefore building a comprehensive new search engine, together with a web-based front end, called Multimedia Commons Search (MMCS). The open-source framework thus far is built around a Solr-based search, connected via Flask to a React web page. It is available at `http://search.mmcommons.org/`, on GitHub, and via Google Drive. The search engine currently uses the YFCC100M metadata; Yahoo-supplied extensions such as geographical information and autotags; video labels [71]; and language labels [79]. Confidence scores for the autotags and video tags can be used to adjust the precision of the search.

Figure 4 shows the current state of the MMCS web interface. The user can generate complex searches (optionally using Solr queries) and use metadata filters to exclude irrelevant results.

We are continuing to add new filters and search types. We next plan to add similarity search (e.g., based on LIRE [89]), using existing and/or new human-generated video labels [e.g., 10], as well as adding some of the existing feature sets to aid in building new filters. We may also generate new autotags

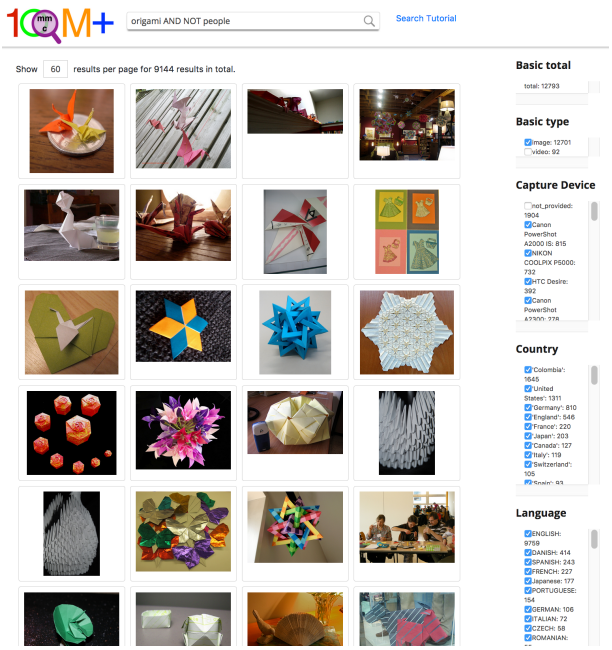

Fig. 4. Screenshot of the Multimedia Commons Search (MMCS) interface.

based on high-frequency user-supplied tags, which can then generalize over the whole dataset [127]. In addition, new autotags could be generated by transferring classifiers trained on other datasets.

Several studies have already investigated how and to what degree the user-supplied tags in the YFCC100M can be used to bootstrap annotations for untagged media. For example, Izadinia et al. [68] suggested a new classifier to deal with noisy class labels. It uses user-supplied tags ("wild tags") to generate automatic (weak) annotations for untagged data in the YFCC100M. Popescu et al. [111] developed an evaluation scheme in which user-supplied tags can be used to evaluate new descriptors (when enough user tags are available)—though again, field studies may require descriptors that do not tend to occur in tags. As we expand further, filters will incorporate any newly generated metadata, such as estimated locations for non-geotagged media. Other additions may involve automatically detecting characteristics likely to be useful across a variety of studies, such as 3D posture specifications. Such estimation combines 2D pose estimation [17, 138] with a mapping from 2D to 3D, mostly based on existing 3D databases of human poses.[2] We also intend to add translation capabilities. And as we noted above, users may choose to contribute new classifiers they train or new annotation sets they create.

Finally, in addition to adding new metadata and new search capabilities, we hope to incorporate additional UGC corpora beyond the YFCC100M, including text and pure audio corpora.

3.3. From Search to Data Processing

To enable scientists to perform studies, developing search filters is not sufficient. A user-friendly data-processing framework is also necessary, for annotating, correcting labels, extracting features from the data, and/or creating new classifiers and filters from search results or from target images they already have on hand.

In addition to improving MMCS, we will add a data processing component, for example by extending the signal processing and classification environment pySPACE [81] to work on multimedia data. To build new classifiers for images and video keyframes, the framework might incorporate CNN features from VGG16 trained on ImageNet [124], train a simple classifier on the examples selected by the user, and then use the classifier to retrieve additional images from the YFCC100M. (Although this piece has not yet been incorporated in the publicly available version of the framework, Section 6 describes a trial run of the process.) As we noted in Sections 3.1 and 3.2, these classifiers and labels could then be incorporated into the framework for future studies, providing more prefab filtering options to new users.

3.4. A Potential Issue: Selection Bias

One issue the data-processing framework will need to help researchers address is selection bias. Bias might arise from the filtering strategies or from the distribution of the dataset itself, potentially affecting the results of a study. Kordopatis-Zilos et al. [80] analyzed several dimensions of bias in the YFCC100M with respect to the location estimation task:

- Location bias: The YFCC100M is biased toward the U.S. and (to a lesser degree) Europe; people are more likely to take pictures in certain places (like tourist destinations);[3]
- User bias: Some users contribute a much higher proportion of the data than others;
- Text description bias: Some data comes with many tags and long descriptions, while some is not even titled;
- Text diversity bias: Some tags and descriptions might be very similar (especially if uploaded together); and
- Visual/audio content bias: Data may contain more or fewer of the particular visual or audio concepts targeted by automatic classifiers.

[2] However, addition of 3D postures may not be feasible for a while yet. In a test of some state-of-the-art pose estimation tools, we realized that current capabilities are more limited than they are often reported to be. The recognition quality was poor when we queried less common (or less straightforward) postures involving, for example, crossed legs or rotated hips. We believe this discrepancy between the reported performance and our results arises from standard issues with transfer from small datasets with a limited set of targets (in this case, poses) to wild UGC data. In addition, in practical terms, pose estimation requires heavy processing, which is rather slow on a dataset the size of the YFCC100M.

[3] In addition, filtering must account for non-unique place names like *Richmond*.

Other important dimensions might include language (of content or metadata), properties of the recording device, time of day, or the gender, age, etc. of the contributors and subjects. For applied studies involving training and evaluating classification algorithms, class imbalance can also be an issue [125].

For MMBDS filtering, we intend to build on the sampling strategy suggested by Kordopatis-Zilos et al. In this approach, the percentage of the difference between a given metric computed on the target dataset compared to a metric computed on a less biased reference dataset is reported ("volatility" in their equation 6). To generate one or more reference datasets, the system can apply strategies that mitigate the aforementioned biases (like text diversity or geographical/user uniform sampling) or separate the biased dataset into several subgroups (like text-based, geographically focused, and ambiguity-based sampling). The search engine can support creation of the reference datasets, and then the data-processing framework can calculate the different performance metrics in an evaluation setting.

In addition to trying to mitigate bias, the data processing framework should make the user aware (e.g., via a visualization) of possible residual biases that could influence their results.

4. EXAMPLE STUDIES: ANSWERING SCIENTIFIC QUESTIONS WITH UGC

A wide range of studies in natural science, social science, and the humanities could be performed or supplemented using UGC media data rather than controlled recording.

In some of the examples in this section, we describe existing studies and suggest how they could be reproduced or extended with the YFCC100M dataset. In other examples, we suggest studies that have not yet been performed at all.

Depending on the example, the UGC data might be the final object of study, or it might act as a pilot. As a pilot or pre-study, it could help researchers get a handle on what variables they most want to examine or isolate in controlled data-gathering, what other variables they need to control for, how much data and how many camera angles they need, etc. It can also alert them to additional factors or possible variables of interest that they might not have expected *a priori*. Such a pilot could save a project significant time and money.

4.1. Environmental Changes and Climate Indexes

Researchers have already begun combining social-media images with data from other sources to analyze changes in the natural world.

In general, it is possible to extract or classify any specific kind of plant, tree, lake, mountain, river, cloud, etc. in images. Since natural scenes are popular motifs in vacation images, there is quite a bit of relevant data in the YFCC100M. This data can be used to analyze changes in those features across time and space. For example, calculated features such as color scores can be used to create indexes, such as a snow index, a pollution index, or an index representing the height of a river or creek.

In one recent case, Castelletti et al. [23] used a combination of traditional and UGC data to optimize a control policy for water management of Lake Como in Italy. They calculated virtual snow indexes from webcam data and from Flickr photos. Using data from even one webcam was already slightly better than using satellite information, and combining the two showed large benefits. However, they were not able to leverage the Flickr photos for similar gains because their image dataset covered too short a time period.

By using the ten years of YFCC100M data with this approach, researchers could further improve on these results, and even generalize to other regions of the world where YFCC100M coverage is dense (especially North America, Europe, the Middle East, and Australia). In related studies [135, 146], satellite images were taken as ground truth to estimate worldwide snow and vegetation coverage using (unfiltered, but geotagged and time-stamped) Flickr images. Using a combination of these two approaches to enhance snow indexes all over the world would be very useful for climate analysis.

To estimate air pollution ($PM_{2.5}$ index) from images, recent approaches have used image data from small, purpose-built datasets. For example, Liu et al. 2016 [87] correlated images with publicly available measurements of the $PM_{2.5}$ index in Beijing, Shanghai, and Phoenix, using specific features to construct estimators of the air pollution based purely on images. Zhang et al. [145] used a CNN-based approach to the same task, focusing on Beijing. Extending those studies to larger datasets (different points of interest, different distances to observed objects, and different times and seasons) and more locales could enable the generation of air pollution estimates for places where no sensors exist. This could be achieved by correlating geotagged, time-stamped outdoor images in the YFCC100M dataset with pollution measurements for locations where that data is publicly available, then translating the results to locations without pollution sensors.

Cloud-cover data is also highly relevant for longterm analysis of the natural world [42]. Some work has been done on automatically classifying cloud types [140]. However, methods for globally complete cloud-cover estimation have not been developed to extend localized automatic detection and human-generated estimates, which often suffer from gaps. Existing studies [e.g., 42] could be augmented by adding YFCC100M image data to existing cloud-cover databases. This image data could be gathered by using image segmentation [30] and/or classifiers to pick out clouds, possibly augmented by incorporating PoseNet to determine the direction of the camera [75].

In the related field of geography, there is already great

interest in using UGC to address scientific research questions (a form of citizen science). For example, crowdsourcing has been used to gather data on forest diseases [35], and Flickr data has been used to improve landcover maps [47].

4.2. Human Language and Gesture Communication

The YFCC100M contains a wealth of data on human interaction and communication [79], which could be quite valuable for linguistics, cognitive science, anthropology, psychology, and other social sciences.

In addition to searching for and filtering relevant videos using text metadata and location, researchers could target specific situations using automatic classification functions like speech/non-speech detection, language identification, emotion/affect recognition, and pose recognition. Language identifications on the YFCC100M metadata [79] have already been added to the MMBDS framework. Off-the-shelf speech/non-speech detection [e.g., 139] and speech recognition [e.g., 61, 83, 112], written [e.g., 88] and spoken [e.g., 64] language identification, and speech-based emotion recognition [e.g., 48, 133] packages could also be incorporated.[4] Extracting 3D models for pose recognition has been studied for images [11, 24, 143], and will hopefully continue to improve and become more efficient. If so, this work can be extended to video [147], in combination with work on motion trajectories already being done with YFCC100M videos [18]. This aspect would be quite challenging, but very useful for several of the examples described in this section and Section 5.

An example of a study that could be expanded in this way compares how people talk to pets, babies, and adults. Mitchell's (2001) analysis identified ways in which people in the U.S. speak to their dogs as they would to infants, in terms of content (short sentences) and acoustic features (higher pitch), and ways the two types of speech differ [99].

With MMBDS, the dataset for this study could be broadened, and comparisons made to how people talk to other pets and non-domestic animals, along with comparing child-directed and animal-directed speech in other cultures. For such a study, short videos like those on Flickr would be sufficient. A large number of videos could be gathered using metadata searches, given the popularity of animals and children as video subjects in YFCC100M; in addition, some videos already have strong annotations for interaction with animals [10]. If necessary, this could be supplemented with speech/non-speech detection and other feature-based filters to identify babies, children, and pets.

On a much wider scale, there are many topics in child language acquisition that could be explored using UGC data, especially for high-frequency phenomena. Existing corpora of children's speech and child-directed speech (CHILDES [90] being the most widely used) usually include some video data, but by far the majority is audio-only or annotated transcripts. This limits researchers' ability to examine the relationship between children's utterances and the situational context for them (for example, what they might be trying to describe or achieve). Acquisition researchers therefore often spend a significant portion of their budget on video recording—and a significant amount of time dealing with Institutional Review Boards' requirements for data involving children.

We describe here two among the many examples where an important acquisition study could be extended using UGC data. In one example, Choi and Bowerman (1991, 2003) examined how English- and Korean-speaking children conceptualized the relationships between two objects [13, 29]. Different languages highlight different aspects of spatial relationships; for example, English *put in* vs. *put on* distinguish containment from surface attachment, while Korean *kkita* vs. *nehta* distinguish close-fitting from loose-fitting relationships. Choi and Bowerman used videotapes of both spontaneous speech and controlled experiments to investigate how the difference in language affected children's spatial reasoning. They and other authors have since extended this work to, e.g., Dutch [12] and Tzotzil [37].

The YFCC100M could be used to collect data for many more languages, using language identification, detectors for children or children's voices, and location metadata, perhaps combined with tag searches and/or object or pose detectors to identify particular target situations. A subfield of language acquisition explores how children's language learning is integrally related to learning about social behavior as a whole [52, 105, 119]; such studies require a large amount of video data to get the necessary rich context. In one seminal study, Ochs and Schieffelin (1984) used data (including video) from several of their past projects to identify important differences in how caregivers in three cultures talk (or don't talk) to prelinguistic infants [104]. These differences stem from varying assumptions about what kinds of communicative intentions an infant could have.

For other researchers who are extending the findings about caregiver assumptions in large comparisons across many culture groups, in-depth data-gathering is of course necessary. But to answer some preliminary questions and ascertain which cultural groups might follow which general patterns in addressing infants—i.e., to decide where to conduct that in-depth data-gathering—a pilot study using short, uncontrolled videos from a UGC corpus with worldwide coverage could be very helpful. In this case, location metadata could be combined with language identification and identification of babies and children in the videos (and/or broad tag searches) to find potential videos of interest. Again, using such a pilot to prepare for more controlled, high-quality data gathering is especially helpful for studies involving children and conducted across national borders, given the added dif-

[4] Of course, no available off-the-shelf detector for any video or image characteristic will provide perfect accuracy. However, used in combination with other filters and given the ability to adjust the desired certainty, they can be a valuable tool for narrowing or broadening the search space.

ficulties in scheduling data-gathering and getting proper permissions.

The related and growing field of gesture studies relies for obvious reasons on video-recorded data [e.g., 34, 106, 137]. Here, again, there are questions that can be answered using short videos or even images. Depending on the question, UGC might provide the dataset for study or act as a pilot. At the least, even uncontrolled, messy UGC data can give the researcher a preliminary sense of how frequent a particular phenomenon is and whether it is common across speakers or across a culture.

But in the case of gesture, tag-based search will likely produce little of value. Gesture researchers interested in systematic description of the ordinary hand movements, postures, and facial expressions that accompany normal conversation will not be able to find those ordinary gestures using tags; after all, tags tend to point out the exceptional.

To take a concrete example, a gesture researcher at one of our institutions wanted to study when people gesture with a pointed finger but without pointing at anything in specific (for example, how does it correlate with emphatic tone of voice?). However, when she tried a tag search for *pointing* in UGC videos, of course, she found either extreme examples (shaking a pointed finger angrily) or examples where people were pointing at something or someplace, rather than the small hand movements she wanted to analyze. In this case, the researcher gave up on pursuing the question—but if she had been able to use feature-based query-by-example, or (better yet) initiate a search by specifying the 3D relationship between the hand and fingers, she could have had much better success.

4.3. Human Behavior: Emotion Examples

One area of behavioral research where multimedia data is vital is the study of how humans express emotion, and how we understand others' emotions. (And human emotion research in turn feeds into multimedia research on automatic emotion understanding; see Section 5.2 for examples.)

In particular, it can be difficult to obtain spontaneous recordings of a wide range of emotions in an experimental setting. Researchers can set up situations to try to elicit emotional reactions (sometimes called "induced emotion") [e.g., 63], but there are limits to this practice—especially given the ethical requirements to obtain permission. UGC therefore has the potential to fill a large gap in emotion and affect research that is only beginning to be addressed.

However, as with the gesture example discussed in Section 4.2, a simple approach to tag-based search—e.g., using emotion words like *disappointment*—is not likely to yield scientifically useful results. It will likely only turn up examples where the behavioral expression of that emotion is extreme, and/or where the uploader has some personal reason for commenting on it. Researchers targeting the expression of specific emotions can find more representative samples by searching for situations likely to elicit those emotions, either using tags, existing event annotations, or detectors for events or other aspects of the situation.[5] For example, sports events are associated with feelings of excitement, suspense, triumph, and disappointment [95].

Alternatively, a researcher could start by searching for particular facial expressions, gestures, or tones of voice—either using feature-based query-by-example or by specifying 3D postures, motion trajectories, pitch contours, etc.—then analyze the types of situations that lead up to those reactions and what, if anything, the participants say about them.

These avenues of multimedia research could be very helpful in developing a fuller picture of the wide range of behaviors that can express any given emotion. Much has been done to identify the most prototypical facial expressions, vocal inflections, etc. associated with particular emotions [43, 44, 60]. However, one of the important questions in the field is how to get beyond those prototypical reactions to a more comprehensive understanding—especially as emotion expression is known to vary quite widely even within a single culture, much less across cultures [45, 115].

The flipside of research into emotion expression is research into how people interpret and categorize the emotions of others based on their behaviors. Image, audio, and video data are of course a mainstay for creating test stimuli in such experiments. However, much of the most prominent research on emotion recognition [e.g., 43, 60] has used acted rather than spontaneous emotion, which (besides being unnatural) tends to stick to prototypical cues. Researchers are discovering the limits of this approach and the questions it can address [e.g., 103] (including for translation into automatic affect recognition; see Section 5.2). Hence, the last few years have seen a shift to recognizing the need for more spontaneous data, such as that found in the YFCC100M. After all, humans can and do deal with quite messy data about the emotions of the people around them [49].

One possible research project (to potentially be conducted by one of the authors) would be to use YFCC100M data to compare how speakers of different languages conceptualize and talk about emotions. For example, there is a large class of languages that express virtually all emotions as states of the experiencer's body parts [94]. This is particularly common in Southeast Asia, as in the following example from the Hakha Lai language of Burma: *ka-ha-thi na-thak*, literally 'my tooth-blood you itch', meaning 'I can't stand you' [129]. Geotagged videos from Southeast Asia could help to investi-

[5] While it is possible that off-the-shelf emotion/affect detectors could be used here, their utility is limited if the object of study is emotion expression or interpretation itself. Automatic classification of human behavior is necessarily always a few steps behind what is known in behavioral science [55, 144]. Creative approaches can at least partially account for this [92], but a simple system can only confidently identify the most prototypical, unambiguous examples, rather than the full range. (In addition, such detectors are often trained on acted emotion—see Section 5.2.)

gate whether other, non-verbal aspects of emotion expression differ in related ways.

Another potential study by the same researcher might investigate how emotion categorizations and descriptions are influenced by an understanding of context. A broadly representative set of test stimuli could be compiled using the methods described above, along with location metadata, existing strong subset annotations, and (ideally) language identification, to find likely candidate images and videos from different languages and cultures. Researchers could then compare how people described the emotions of the people in them depending on whether they were shown the preceding (or surrounding) context, or only the snippet of the person expressing the target emotion. An extended experiment could also compare descriptions based on only the audio or only the visual stream from a video.

4.4. Location-Based Comparisons

In the YFCC100M, around 50% of the data is geotagged, and location estimations can be generated for many of the remaining images and videos, especially ones recorded outdoors [26, 28]. From this, it is possible to approximately infer the location of many users' homes—or at least their hometowns—as well as where they travel to. Even where location estimation for a given image is too difficult (especially indoors), information can be gleaned at the user level based on the individual's other uploaded images. For some studies, it may even be more helpful to know where the user is from than where the picture was taken.

One possible research question in this area would be to compare where people from different locales like to travel (and take pictures) [32, 72]—do they go to (other) urban areas? Do they go out into nature? At what times of year? A more in-depth analysis could examine changes in those behaviors over time to identify how preferred tourist spots change in response to world events.

As another example, a researcher (e.g., in anthropology or marketing) could use geotagged indoor images and videos to identify patterns in the personal possessions of people from different locations and backgrounds. Using classifiers on labeled data, it would be possible to determine the brands and values of at least some items. In many cases, it might be possible to identify the objects and their values by image comparisons, e.g., with an online store. Companies could also gather data about how their products are used in practice in different cultures and countries, and use this information to develop new services and products or marketing strategies [126]. Such analysis could lead to a variety of automated applications as well; see Section 5.4.

As we noted in Sections 4.2 and 4.3, location data can also be used in cross-cultural studies. As a starting point, styles of photography could themselves be compared across locations and over time.

As another example, comparisons of gender presentation and gender dynamics across cultures are usually based on in-depth fieldwork on the ground. But such studies could be supplemented with UGC data to provide cross-checking against many more data points for a given culture, and to quickly gather at least some data from many different locales without having to travel to all of them. A researcher could identify such data using geotags and (optionally) inferred locations, relevant user-supplied tags, person detectors, and possibly language detectors.[6]

4.5. Medical Studies

Wang et al. (2017) used a combination of machine-learning methods to attempt to identify Flickr users who engage in deliberate self-harm [136]. They showed differences by text characteristics, user profile statistics, activity patterns, and image features. Classification results were not as accurate as is usual for more well-studied tasks, but were certainly accurate enough to produce a good candidate set for a field expert to narrow down. Wang et al. suggest that data gathered via such a detector could help researchers enrich their understanding of the triggers and risk factors for self-harm, along with studying the self-presentation and interactions of self-harmers on social media per se.

That work had a specific topical focus (on content that might not be well represented in Creative Commons media), and thus approached the problem somewhat differently than we are proposing for a more multi-purpose search interface. However, we consider the results to be a promising indicator of the potential of such efforts.

We have not tried to quantitatively assess how much content can be found in the YFCC100M to represent abnormal or pathological behavior or physical conditions. However, even for cases where the YFCC100M does not contain many examples of a targeted condition, it can still be quite useful to researchers studying that condition: It can provide a quick and easy way to gather a baseline dataset to compare to. (In fact, Wang et al. pulled their examples of non–self harm content from the YFCC100M, though they did not target any specific behaviors for that control set [136].)

5. EXAMPLE UGC-BASED AI APPLICATIONS

As with the examples in Section 4, some of the possibilities we describe for applied research with UGC extend existing studies, while some of these areas have not been explored much at all.

[6] A potential filter could be built to exclude tourists' contributions where they are unlikely to be apropos, using tags and inference from the locations of other pictures from that user.

5.1. Movement Training for Robotics

Imitation learning in robotics requires recorded data of humans performing a target behavior or motion. Movements can be learned and transferred via shadowing [e.g., 4, 98], for example to a robot arm for grasping [142] or throwing, or an agent can learn from recordings of a single human over a long time-span [69].

Taking throwing as an example, it would be possible to infer joint and object positions, velocities, and accelerations from movements observed in data from a variety of sources, including targeted motion-tracking systems but also recordings of, e.g., ball games. Within our framework, 3D human pose estimation would have to be applied to single keyframe images from a video [11, 24, 143, 147].

After extracting the movements, the data can be segmented, for example using velocity-based multiple change-point inference [122]. The motion primitives can then be post-processed and classified [59]. Finally, the movement behavior can be optimized and transferred [58].

Such movement trajectories can be used not only for transfer learning, but also for more general analyses of movement patterns and to build classifiers, for example to identify movement disorders.

5.2. Interaction Training for AI Systems

As robots, dialogue systems, AI assistants, and other AI-based services improve in sophistication and interactivity, they need to be able to recognize and categorize not just speech, but human emotions, attentional cues, and other cues that can help in interpreting intent.[7] There is therefore a major drive (for example, in affective computing) toward automatic recognition of emotion[8] in multimedia data, including facial expressions [surveys in 56, 107, 117], gesture/posture [survey in 76], vocal cues [surveys in 14, 134], and biological signals [surveys in 5, 56]—and combinations of those modes [surveys in 19, 39, 144].

However, as we noted in Section 4.3, it can be difficult to get truly spontaneous, naturalistic data for emotion expression—to use as training data for automatic systems, as much as for the scientific purposes mentioned above [19, 134, 144].

For this reason, automatic affect recognition researchers have often used datasets of acted emotion [e.g., 7, 74, 109]. While the situation has improved in recent years, with several annotated video datasets of "spontaneous" emotion expression being released, much of that data has in fact been induced emotion, collected under contrived conditions [e.g., 96, 109, 113], or at best from television interviews [e.g., 40, 54]. Even "spontaneous" datasets are usually collected under contrived conditions, either induced or (at best) interviews. (In addition, such data are often collected under ideal conditions, in terms of lighting, head angle, etc.) Available annotated datasets collected in more truly naturalistic situations with rich context are fewer and are often audio-only [e.g., 15, 62] (where it is easier to minimize the effects of recording [38]).[9] The value of UGC data for this purpose is therefore coming to be recognized with new datasets [e.g., 8], but none as yet have strong human-generated annotations.

Comparisons show that acted and induced or spontaneous emotion expression can differ in ways that have consequences for recognition [e.g., 33, 120]—fundamental enough that different features may be more discriminatory [e.g., 16, 118, 132]. Most tellingly, cross-testing between datasets shows that training systems to recognize acted, prototypical, or larger-than-life examples of affective cues does not necessarily prepare them well for their actual task: recognizing what humans are doing "in the wild" [86, 101, 121, 128] (though model adaptation can help [1])—nor vice versa, for that matter [131].[10] Natural emotion expression may be much subtler, and cues may be ambiguous—either because there are multiple emotions that cue commonly expresses or because the person is expressing an emotional state that does not fall neatly into one category [e.g., 36, 41, 50, 56, 77, 101]. In addition, even dimensional (non-categorical) analysis systems run into problems because of the wide variation between individuals [16, 53, 57, 108].

To take just one example of an intriguing study that invites replication with naturalistic data, Metallinou et al. (2012) were able to improve emotion recognition in video by taking into account prior and following emotion estimations, i.e., by modeling the emotional structure of an interaction at the

[7] However, it is worth pointing out that UGC data also has the potential to aid in automatic speech recognition (ASR) in arbitrarily noisy environments, especially in recognizing children's speech. Again because of the additional necessary permissions and precautions, ASR researchers have a much smaller pool of recorded corpus data to draw from for recognizing the speech of children, in comparison to the available resources for adult speech [31].

[8] We use *emotion (expression)* and *affect* somewhat interchangeably in talking about the whole field; for particular applications, researchers often draw finer distinctions.

[9] By "truly naturalistic", we mean occurring naturally in the course of everyday life, rather than induced by a researcher. We do not (necessarily) mean what is sometimes called in behavioral research "biologically driven" (or, even more confusingly, "spontaneous") emotion expression [e.g., 91], which is then opposed to some general category of unnatural or non-spontaneous behavior that lumps together learned or socially driven emotion expression with acted or deliberately false emotion.
UGC image and video data will include a large proportion of emotion expression driven by communicative purposes and manifested according to learned cultural conventions (and varying in how closely it represents what the person is "really" feeling), but it is nonetheless a response to a naturally arising context [95, 116]. For an AI system, it will be important to be able to interpret (and emulate [46]) both cues that can be consciously mediated (such as vocal pitch) and cues that (largely) cannot (such as pulse rate)—and to use the two different types of information appropriately in interaction. (Here, we set aside the effects of being recorded in the first place, which are nearly unminimizable for any kind of video data [38].)

[10] Many of these comparisons are between acted and induced data. Extending to three-way comparisons with UGC data (of comparable quality) is one obvious starting point for research.

same time as the individual cues [97]. However, they used dialogues improvised by actors rather than naturally occurring emotional interactions.

An important question is whether they would have achieved similar results using non-acted data. In other words, is automatic recognition of emotions in the wild helped in the same way or to the same degree by taking into account prior and following judgments, or is there something special about the unity of presentation in an acted situation? To investigate this, researchers could collect and annotate data from the YFCC100M videos that depict interactional situations similar to those improvised by the actors, using tag search, speech detectors, and recognizers for particular activities. For example, the existing autotags and video event labels could be used to filter for targets like *people*, *sport*, *wedding*, *hospital*, *fight*, or *love*.

For humans, context and world knowledge can help to disambiguate or clarify confusing emotional cues [49], but context-sensitive automatic affect interpretation is still a fairly young field [e.g., 21, 36, 85, 97]. As Castellano et al. 2015 (among others) point out, AI systems require both training and testing data specific to the types of interactional situations they are likely to face, to ensure they will interpret cues correctly and react appropriately [19]. They thus point out the need in affective computing for more datasets drawn from recordings of specific natural contexts.

TheThe YFCC100M offers an opportunity to gather corpora of motion expression occurring in natural situations and recorded under non-ideal conditions, with a wide variety of contexts and drawn from a variety of cultures. These corpora could then be annotated and used to train systems to interpret emotion relative to situational and cultural context, and to produce interactional styles matched to those contexts and to the culture [93]. The fact that a given Flickr user will often have videos of the same people (such as the user's family members) across many different situations can also provide an opportunity to isolate individual variation from contextual variation.

To demonstrate the importance of context-specificity, Castellano and colleagues collected data from children playing chess—in this case, with a robot—to train the robot to interact socially while playing the game [20, 22]. They found that including context features in the affect recognition component increased children's engagement with the robot, and that game-specific context features had a bigger effect than general social context features.

However, playing chess with children is obviously only one of thousands or millions of situations AIs might be called upon to interact in, each with its own norms and scenarios that might be broken down into context features. Whether one is building many purpose-specific AIs or a more general-purpose AI system, it would quickly become prohibitively expensive to try to collect a similarly controlled emotion dataset for every situation.

A researcher could use the MMBDS framework to gather video data on situations of interest, combining tag and location search with feature-based query-by-example, existing strong labels (e.g., event labels), and recognizers for, e.g., human faces, speech/non-speech, and particular relevant objects. For common situations such as tourist interactions, researchers may be able to find enough high-quality videos from useful angles to constitute a training dataset in itself, or at least enough to identify which context features to target (i.e., to develop a codebook). For other situations, YFCC100M data could help researchers prioritize what kind of controlled data to collect, identify the range of conditions they might need to get a representative sample, and find alternatives or proxies for situations where controlling data collection may be too difficult.

5.3. Construction Site Navigation

For autonomous driving systems, construction sites are very challenging [6]. Each site looks different, with large differences between countries. In addition, it is difficult to get sufficient data. For relatively predictable features of sites such as warning signs, existing classifiers [123] could be applied or extended relatively easily. In the MMBDS framework, more sophisticated approaches to data-gathering could be implemented to address more challenging and heterogeneous characteristics.

Using archived traffic reports or government data to locate construction sites, a researcher could then filter out potentially relevant videos and images from the YFCC100M dataset using timestamps and geotags. This data could be quickly reviewed to verify its relevance and perhaps add additional tags. This approach would extract a much bigger dataset than could be obtained solely using warning-sign classifiers. The next step would be to train classifiers and segmentation algorithms to recognize and characterize construction sites (even if they are unmarked), and to assess the likelihood of various complications and risks.

As a side note, this is a good example of a case where maximal applicability requires having human-interpretable algorithms, so the knowledge gained can be best integrated into autonomous driving applications.

5.4. Other Location-Based Applications

In addition to studies of how people differ by location and cultural background, such as those we suggested in Section 4.4, the YFCC100M location information can be used in combination with other data extracted from images/videos for several AI applications. For example, a number of studies have looked at using social-media content—including YFCC100M data—to automatically generate tourist guides [25, 78].

The studies of people's possessions proposed in Section 4.4 could be extended to automatic applications, for example for targeted advertising and market research. Data on

possessions could also be used to generate features and train an estimator for housing values based on public property-value data, then to transfer this classification to regions that do not have publicly available data on housing values.

Conversely, the dangers posed by multimedia analysis techniques like automatic valuation and location estimation are an important topic in online privacy, where researchers are examining and measuring how much private information users give away when they post, e.g., images and videos [27]. The potential for new applications like location-aware classification of people's possessions increases those dangers. After all, such techniques—especially when combined with information from other online sources—could be used not only by marketers but by criminals, for example to plan a robbery (called *cybercasing* [51]).

6. IMPLEMENTATION CASE STUDY: ORIGAMI

This section describes a practical case study we conducted to analyze the requirements for the MMBDS framework. To our knowledge, this is the first time that machine learning has been used in the study of origami.

6.1. Background: Origami in Science

Origami is the art of paperfolding. The term arises specifically from the Japanese tradition, but that tradition has been practiced around the world for several centuries now [65]. In addition to being a recreational art form, in recent years, origami has increasingly often been incorporated into the work and research of mathematicians and engineers.

For example, in the 1970s, to solve the problem of packing large, flat membrane structures to be sent into space, Koryo Miura used origami: he developed a method for collapsing a large flat sheet to a much smaller area so that collapsing or expanding the sheet would only require pushing or pulling at the opposite corners of the sheet [100]. Similarly, Robert Lang is using computational origami-based research to help the Lawrence Livermore National Laboratory design a space telescope lens ("Eyeglass") that can collapse down from 100 meters in diameter to a size that will fit in a rocket roughly 4 meters in diameter [67, 84]. In the field of mathematics, Thomas Hull is investigating enumeration of the valid ways to fold along a crease pattern (i.e., a diagram containing all the creases needed to create a model) such that it will lie flat. He uses approaches from coloring in graph theory to solve the problem [66]. In medicine, Kuribayashi et al. used origami to design a metallic heart stent that can easily be threaded through an artery before expanding where it needs to be deployed [82].

In other words, origami is becoming an established part of science. To support research on origami, we decided to generate a large origami dataset, building around data in the YFCC100M.

6.2. Potential Research Question: Regional Differences

As our test case for the requirements analysis, we gathered a dataset from the YFCC100M that could be used to answer questions about regional variation, such as: What are the differences between countries/regions in terms of what styles and subject matter are most popular? How do those differences interact with other paperfolding traditions?

Some traditional approaches to this question might be to look at books or informational websites about origami from different places, or to contact and interview experts in those well-known places. However, the books and websites are unlikely to be comprehensive across regions, and experts might not know about (or be concerned about) origami in all the places it is practiced. Alternatively, one could travel the world, visiting local communities to gather data about origami practices, but this would be very expensive and time-consuming (if one could even get funded to do it).

However, origami can also fruitfully be studied using UGC media. People are often proud of their origami art, especially when it is of high difficulty and quality. It is therefore the kind of thing that people take pictures of, and upload them to social media like Flickr.

Using the MMBDS framework to target location-specific images of origami, data for such a field study could be gathered in a day.

6.3. The Limitations of Text-Based Search

We began by assessing what could be gathered via a simple text metadata search, using the YFCC100M browser [73]. The keyword *origami* netted more than $13,000$ hits. However, only about half of the returned images were geotagged, and we identified several issues with the remainder.

Most obviously, more than 30% of the images did not contain origami. In addition, the spatial distribution map did not match where common sense tells us origami should be prevalent. The most uploads came from Colombia, followed by the U.S. and Germany. Japan was in seventh place, and there were only two examples from China.

This unexpected distribution likely had several causes. First, there is a general bias towards the U.S. in the YFCC100M [73]. It was gathered from Flickr, which is most popular in the U.S. and Europe, and less popular elsewhere (and in the case of China, was blocked for part of the target period). Second, subset geographical skew can also stem from user bias; in this case, nearly $1,000$ images were uploaded by one artist (Jorge Jamarillo) from Colombia. Finally, a search on *origami* would not catch examples tagged in Japanese characters. Media whose metadata is in a different language than the researcher is searching in will not be included, putting the burden of language-guessing and translation on the researcher. (Though many Flickr users do include English tags, whatever other languages they use [79].) Multilingual search is also limited by character encoding issues;

at present, a researcher could not search using Japanese at all. This highlights the general problem described in Section 3, that searching the text metadata will miss many examples, most obviously those that do not have text metadata at all.

These limitations show that we need a more comprehensive search engine that can consider multimedia content. As we described in Section 3.2, the user should be able to create new filters by selecting good examples. For a location-based study like this one, location estimation could expand the dataset. Furthermore, the search needs to incorporate translation, either of the search terms or the media metadata. Finally, a science-ready search engine should allow a researcher to quantify and visualize bias (such as user bias) and have a configurable filtering tool for reducing bias (see Section 3.4).

6.4. Data Processing and Filter Generation

For an effective search, statistics over the metadata are required to quickly generate ideas for additional search terms to include or exclude. In our case, we narrowed our search by using the prominent tags *papiroflexia* (Spanish for *origami*) and *origamiforum*, and found they were more reliable than *origami*. We used these terms, plus minimal hand-cleaning, to collect $1,938$ geotagged origami images for the first part of our filter-training dataset.[11]

That process highlighted some considerations for the selection and data-processing framework (as described in Section 3.3)—and why it needs to be part of the same system. For example, to create a training dataset for filter generation, we wanted to use images where the origami object was dominant (rather than, for instance, a person holding an origami object). Instead of having to hand-prune the initial results, it is much faster to begin with automatic annotations for, e.g., people (such as the existing autotags, which we added later), or with similarity-based filters.

We began with our extracted dataset of $1,938$ examples to analyze the requirements for the process of generating new filters. First we applied a VGG16 neural network [124], trained on ImageNet with $1,000$ common classes (not including *origami*). The top-1 predictions were spread across 263 different classes, and the top-5 predictions were spread across 529 classes.

The most common classes we found are summarized in Table 1. Our ground truth origami images were quite often classified as *pinwheel*, *envelope*, *carton*, *paper towel*, *packet*, or *handkerchief*—all visually (and conceptually) similar to origami/paperfolding in involving paper. We also noted that (beyond *pinwheel*) some images were classified as containing the real-world objects the origami was supposed to represent, such as bugs or candles—and in some cases, like flowers, the origami was often difficult to distinguish from the real object even for a human.

Table 1. Top-i predictions for our extracted YFCC100M subset (n = number of occurrences).

Top-1		Top-5	
name	n	name	n
pinwheel	243	envelope	788
envelope	243	pinwheel	518
carton	117	carton	499
paper towel	76	packet	328
honeycomb	71	handkerchief	314
lampshade	41	paper towel	302
rubber eraser	39	rubber eraser	238
handkerchief	37	candle	218
pencil sharpener	35	lampshade	192
shower cap	34	wall clock	168

We therefore believe that assigning ImageNet labels to the whole YFCC100M dataset via VGG16 could be a first step in improving the search function by allowing multiple filter types. For example, searching for data tagged as *envelope* by the VGG16 net and *origami* in the text metadata would probably deliver cleaner results than just searching for *origami*.

However, since VGG16 trained on ImageNet apparently includes many classes that are at least visually similar to origami, VGG16 features (not classes) would seem to be the better basis for constructing a *new* classifier/filter for origami. In general, features from deep learning networks are quite powerful for image classification [e.g., 102], so are good candidates for use in our framework.

As we pointed out in Section 3.3, if a researcher already has some target images on hand, they can be used to improve the filter. In this case, we added additional data, by scraping images from two origami-specific databases [2, 130]. After some minor cleaning by hand (to remove instructions and placeholder images), these databases yielded $3,934$ and $2,140$ additional origami images. To construct a non-origami class, we used the ILSVRC2011 validation data [114].[12] We used the first 8 examples for each label, excluding *pinwheel*, *envelope*, *carton*, *paper towel*, *packet*, and *handkerchief*. In all, we had $8,011$ images with origami and $7,976$ images without origami.

For features, we used the VGG16 output before the last layer. VGG16 features are already available for the YFCC100M as part of the Multimedia Commons [111]; such precomputed features are essential to speed up processing. We generated features for the rest of the data using MXNet.[13] We evaluated the classifier using the pySPACE framework [81], with a logistic regression implemented in scikit-learn [110] (default settings) with 5-fold

[11] Because of our choice of terms, around 25% of this training dataset consisted of Colombian images. Again, ideally, a search engine should include sampling tools to easily ensure that the resulting detector was not skewed towards a specific country.

[12] We could not use the regular ImageNet data for the non-origami class because that is what the VGG16 net was trained on.

[13] `http://mxnet.io`

cross-validation and 5 repetitions. The classifier achieved a $97.4\% \pm 0.3$ balanced accuracy (BA) [125].

6.5. Using the New Filter to Gather More Data

Applying a new trained neural network to all the YFCC100M images would require a tremendous amount of processing. On the other hand, our approach—using a simple classifier model and precomputed features—enabled the transfer on a very simple computing instance without a GPU.

In total, our classifier identified $1,960,303$ images as *origami*. The histogram distribution of the classification probability scores was $[86.9, 5.3, 3.1, 1.5, 1.1, 0.8, 0.7, 0.3, 0.2, 0.1]$ (as percentages).

Visual inspection of the highest ranked 87 origami images (scores > 0.99999) showed only 2 incorrect identifications. But looking at the lower-scoring images, we found much worse performance. Even for classification probabilities between 0.99 and 0.9, visual inspection of a subset revealed that less than 50% of the images contained origami.

This discrepancy shows that it is crucial to allow the user to adjust the decision boundary for a given filter. It also shows that, at this scale, filtering will generally need to be above 99% BA in quality. Having an error of 1% for the non-target class can easily result in millions of misclassifications, effectively swamping a low number of relevant examples.

Visual inspection also showed that many images were not photos; this suggests that a pre-supplied photo/non-photo filter would allow for quick weeding. This could be done using EXIF data (released as a YFCC100M extension) to select images that have camera information. In addition, once the new *origami* filter had been applied, autotags (or other annotations) could be used to remove types of images that were commonly misclassified as origami (in this case, people, animals, vehicles, and food)

6.6. From Images to Field Study

For this case study, we extracted only those origami/paperfolding images with location information (39% of those found).

To compare regional variations in origami styles and subjects, the researcher would want to begin by dividing the geotagged data into geographic units. For this, we used the YFCC100M Places extension, which has place names based on the GPS coordinates.

We looked at the country distribution of the top $5,167$ images (those scoring higher than 0.99). In total, this high-scoring dataset included images from 178 countries, with 93 of those countries having at least 10 examples for an origami researcher to work with. The top 12 countries all had more than 100 images each.

7. OUTLOOK AND CALL TO ACTION

In this paper, we introduced a cross-disciplinary framework for multimedia big data studies (MMBDS) and gave a number of motivating examples of past or potential real-world field studies that could be conducted, replicated, piloted, or extended cheaply and easily with user-generated multimedia content. We also described Multimedia Commons Search, the first open-source search for the YFCC100M and the Multimedia Commons. We encourage researchers to add their contributions to make the framework even more powerful.

Scientists (including some of the authors of this paper) are integral to building a resource like this. Our discussions about the research topics described in Sections 4 and 5 indicate that there is a high level of interest in having an MMBDS framework like the one we are developing. These discussions are already informing the design, and we will continue to involve these and other scientists to ensure maximum utility and usability. We view these kinds of discussions as essential to shifting the focus of the field from potential impact to actual impact. We encourage more multimedia scientists to get in contact with scientists from other disciplines—from environmental science to linguistics to robotics—and vice versa, to build new on-the-ground MMBDS collaborations.

Acknowledgments

We would like to thank Per Pascal Grube for his assistance with the AWS search server and Rick Jaffe and John Lowe for helping us make the Solr search engine open source. We would also like to thank Angjoo Kanazawa for sharing expertise on 3D pose estimation, Roman Fedorov on snow index detection, David M. Romps on cloud coverage estimation, and Elise Stickles on language and human behavior. Thanks also to Alan Woodley and anonymous reviewers for providing feedback on the paper. Finally, we thank all the people who are providing the datasets and annotations being integrated into our MMBDS framework.

This work was supported by a fellowship from the FITweltweit program of the German Academic Exchange Service (DAAD), by the Undergraduate Research Apprenticeship Program (URAP) at University of California, Berkeley, by grants from the U.S. National Science Foundation (1251276 and 1629990), and by a collaborative Laboratory Directed Research & Development grant led by Lawrence Livermore National Laboratory (U.S. Dept. of Energy contract DE-AC52-07NA27344). (Findings and conclusions are those of the authors, and do not necessarily represent the views of the funders.)

References

[1] Mohammed Abdelwahab and Carlos Busso. 2015. Supervised domain adaptation for emotion recognition from speech. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 5058–5062. DOI:http://dx.doi.org/10.1109/ICASSP.2015.7178934

[2] Gilad Aharoni. 2017. http://www.giladorigami.com/. (2017).

[3] Giuseppe Amato, Fabrizio Falchi, Claudio Gennaro, and Fausto Rabitti. 2016. YFCC100M HybridNet Fc6 Deep Features for Content-Based Image Retrieval. In *Proceedings of the 2016 ACM Workshop on the Multimedia COMMONS (MMCommons '16)*. ACM, New York, NY, USA, 11–18. DOI:http://dx.doi.org/10.1145/2983554.2983557

[4] Brenna D. Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57, 5 (May 2009), 469–483. DOI:http://dx.doi.org/10.1016/j.robot.2008.10.024

[5] Jorge Arroyo-Palacios and D.M. Romano. 2008. Towards a standardization in the use of physiological signals for affective recognition systems. In *Proceedings of Measuring Behavior 2008*.

[6] Vahid Balali, Arash Jahangiri, and Sahar Ghanipoor Machiani. 2017. Multi-class US traffic signs 3D recognition and localization via image-based point cloud model using color candidate extraction and texture-based recognition. *Advanced Engineering Informatics* 32 (2017), 263–274. DOI:http://dx.doi.org/https://doi.org/10.1016/j.aei.2017.03.006

[7] Tanja Bänziger, Marcello Mortillaro, and Klaus R. Scherer. 2012. Introducing the Geneva Multimodal expression corpus for experimental research on emotion perception. *Emotion* 12, 5 (October 2012), 1161–1179.

[8] C. Fabian Benitez-Quiroz, Ramprakash Srinivasan, and Aleix M. Martinez. 2016. EmotioNet: An Accurate, Real-Time Algorithm for the Automatic Annotation of a Million Facial Expressions in the Wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5562–5570.

[9] Julia Bernd, Damian Borth, Carmen Carrano, Jaeyoung Choi, Benjamin Elizalde, Gerald Friedland, Luke Gottlieb, Karl Ni, Roger Pearce, Doug Poland, Khalid Ashraf, David A. Shamma, and Bart Thomee. 2015. Kickstarting the Commons: The YFCC100M and the YLI Corpora. In *Proceedings of the 2015 Workshop on Community-Organized Multimodal Mining: Opportunities for Novel Solutions (MMCommons '15)*. ACM, 1–6. DOI:http://dx.doi.org/10.1145/2814815.2816986

[10] Julia Bernd, Damian Borth, Benjamin Elizalde, Gerald Friedland, Heather Gallagher, Luke Gottlieb, Adam Janin, Sara Karabashlieva, Jocelyn Takahashi, and Jennifer Won. 2015. *The YLI-MED Corpus: Characteristics, Procedures, and Plans*. Technical Report TR-15-001. International Computer Science Institute. arXiv:1503.04250 http://arxiv.org/abs/1503.04250 arXiv:1503.04250.

[11] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. 2016. Keep it SMPL : Automatic Estimation of 3D Human Pose and Shape from a Single Image. In *Computer Vision – ECCV 2016*. Springer International Publishing, 34–36. DOI:http://dx.doi.org/10.1007/978-3-319-46454-1_34 arXiv:1607.08128

[12] Melissa Bowerman. 1996. Learning how to structure space for language: A crosslinguistic perspective. In *Language and Space*, Paul Bloom (Ed.). MIT Press, 385–436.

[13] Melissa Bowerman and Soonja Choi. 2003. Space under construction: Language-specific spatial categorization in first language acquisition. In *Language in Mind: Advances in the Study of Language and Thought*, Dedre Gentner and Susan Goldin-Meadow (Eds.). MIT Press, Cambridge, MA, 387–428.

[14] Carlos Busso, Murtaza Bulut, and Shrikanth Narayanan. 2013. Toward effective automatic recognition systems of emotion in speech. In *Social Emotions in Nature and Artifact: Emotions in human and human-computer interaction*, Jonathan Gratch and Stacy Marsella (Eds.). Oxford University Press, 110–127.

[15] Nick Campbell. 2014. Databases of Expressive Speech. *Journal of Chinese Language and Computing* 14, 4 (2014).

[16] Houwei Cao, Ragini Verma, and Ani Nenkova. 2015. Speaker-sensitive emotion recognition via ranking: Studies on acted and spontaneous speech. *Computer Speech & Language* 29, 1 (2015), 186–202. DOI:http://dx.doi.org/https://doi.org/10.1016/j.csl.2014.01.003

[17] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2016. Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. (Nov 2016). arXiv:1611.08050 http://arxiv.org/abs/1611.08050

[18] Carmen Carrano and Doug Poland. 2015. Ad Hoc Video Query and Retrieval Using Multi-Modal Feature Graphs. In *Proceedings of the 19th Annual Signal & Image Sciences Workshop (CASIS Workshop)*. https://casis.llnl.gov/content/pages/casis-2015/docs/poster/Carrano_CASIS_2015.pdf

[19] Ginevra Castellano, Hatice Gunes, Christopher Peters, and Björn Schuller. 2015. Multimodal Affect Recognition for Naturalistic Human-Computer and Human-Robot Interactions. In *The Oxford Handbook of Affective Computing*, Rafael Calvo, Sidney D'Mello, Jonathan Gratch, and Arvid Kappas (Eds.). Oxford University Press, New York, 246–257. DOI:http://dx.doi.org/10.1093/oxfordhb/9780199942237.013.026

[20] Ginevra Castellano, Iolanda Leite, and Ana Paiva. 2017. Detecting perceived quality of interaction with a robot using contextual features. *Autonomous Robots* 41, 5 (2017), 1245–1261. DOI:http://dx.doi.org/10.1007/s10514-016-9592-y

[21] Ginevra Castellano, Iolanda Leite, André Pereira, Carlos Martinho, Ana Paiva, and Peter W. McOwan. 2012. Detecting Engagement in HRI: An Exploration of Social and Task-Based Context. In *2012 International Conference on Privacy, Security, Risk, and Trust and 2012 International Confernece on Social Computing*. 421–428. DOI:http://dx.doi.org/10.1109/SocialCom-PASSAT.2012.51

[22] Ginevra Castellano, Iolanda Leite, Andre Pereira, Carlos Martinho, Ana Paiva, and Peter W. McOwan. 2013. Multimodal affect modelling and recognition for empathic robot companions. *International Journal of Humanoid Robotics* 10, 1 (2013), 1–23.

[23] Andrea Castelletti, Roman Fedorov, Piero Fraternali, and Matteo Giuliani. 2016. Multimedia on the Mountaintop: Using Public Snow Images to Improve Water Systems Operation. In *Proceedings of the 2016 ACM Conference on Multimedia - MM '16*. ACM Press, New York, New York, USA, 948–957. DOI:http://dx.doi.org/10.1145/2964284.2976759

[24] Ching-Hang Chen and Deva Ramanan. 2017. 3D Human Pose Estimation = 2D Pose Estimation + Matching. (Dec 2017). arXiv:1612.06524 http://arxiv.org/abs/1612.06524

[25] Yan-Ying Chen, An-Jung Cheng, and Winston H. Hsu. 2013. Travel Recommendation by Mining People Attributes and Travel Group Types From Community-Contributed Photos. *IEEE Transactions on Multimedia* 15, 6 (Oct 2013), 1283–1295. DOI:http://dx.doi.org/10.1109/TMM.2013.2265077

[26] Jaeyoung Choi and Gerald Friedland (Eds.). 2015. *Multimodal Location Estimation of Videos and Images*. Springer International Publishing, Cham. DOI:http://dx.doi.org/10.1007/978-3-319-09861-6

[27] Jaeyoung Choi, Martha Larson, Xinchao Li, Kevin Li, Gerald Friedland, and Alan Hanjalic. 2017. The Geo-Privacy Bonus of Popular Photo Enhancements. In *Proceedings of the 2017 ACM International Conference on Multimedia Retrieval (ICMR '17)*. ACM, New York, NY, USA, 84–92. DOI:http://dx.doi.org/10.1145/3078971.3080543

[28] Jaeyoung Choi, Roger Pearce, Doug Poland, Bart Thomee, Gerald Friedland, Liangliang Cao, Karl Ni, Damian Borth, Benjamin Elizalde, Luke Gottlieb, and Carmen Carrano. 2014. The Placing Task: A Large-Scale Geo-Estimation Challenge for Social-Media Videos and Images. In *Proceedings of the 3rd ACM Multimedia Workshop on Geotagging and Its Applications in Multimedia (GeoMM'14)*. ACM Press, New York, 27–31. DOI:http://dx.doi.org/10.1145/2661118.2661125

[29] Soonja Choi and Melissa Bowerman. 1991. Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition* 41 (1991), 83–121.

[30] Mircea Cimpoi, Subhransu Maji, and Andrea Vedaldi. 2015. Deep filter banks for texture recognition and segmentation. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 3828–3836. DOI:http://dx.doi.org/10.1109/CVPR.2015.7299007

[31] Felix Claus, Hamurabi Gamboa Rosales, Rico Petrick, Horst-Udo Hain, and Rüdiger Hoffmann. 2013. A survey about databases of children's speech. In *Proceedings of INTERSPEECH*. 2410–2414.

[32] Maarten Clements, Pavel Serdyukov, Arjen P. de Vries, and Marcel J.T. Reinders. 2010. Using Flickr geotags to predict user travel behaviour. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '10)*. ACM Press, New York, New York, USA, 851. DOI:http://dx.doi.org/10.1145/1835449.1835648

[33] Jeffrey Cohn and Karen Schmidt. 2004. The timing of facial motion in posed and spontaneous smiles. *International Journal of Wavelets, Multiresolution,*

and Information Processing 2, 2 (March 2004), 1–12. http://www.worldscientific.com/doi/abs/10.1142/S021969130400041X?journalCode=ijwmip

[34] Eliza L. Congdon, Miriam A. Novack, and Susan Goldin-Meadow. 2016. Gesture in Experimental Studies. *Organizational Research Methods* (2016), 1094428116654548. DOI:http://dx.doi.org/10.1177/1094428116654548 arXiv:http://dx.doi.org/10.1177/1094428116654548

[35] John Patrick Connors, Shufei Lei, and Maggi Kelly. 2012. Citizen Science in the Age of Neogeography: Utilizing Volunteered Geographic Information for Environmental Monitoring. *Annals of the Association of American Geographers* 102, 6 (Nov 2012), 1267–1289. DOI:http://dx.doi.org/10.1080/00045608.2011.627058

[36] Roddy Cowie, Ellen Douglas-Cowie, and Cate Cox. 2005. Beyond emotion archetypes: Databases for emotion modelling using neural networks. *Neural Networks* 18, 4 (2005), 371–388. DOI:http://dx.doi.org/https://doi.org/10.1016/j.neunet.2005.03.002

[37] Lourdes de León. 2001. Finding the richest path: Language and cognition in the acquisition of verticality in Tzotzil (Mayan). In *Language Acquisition and Conceptual Development*, Melissa Bowerman and Stephen Levinson (Eds.). Cambridge University Press, 544–565. DOI:http://dx.doi.org/10.1017/CBO9780511620669.020

[38] Laurence Devillers, Laurence Vidrascu, and Lori Lamel. 2005. Challenges in real-life emotion annotation and machine learning based detection. *Neural Networks* 18, 4 (2005), 407–422. DOI:http://dx.doi.org/10.1016/j.neunet.2005.03.007 Emotion and Brain.

[39] Sidney D'Mello and Jacqueline Kory. 2012. Consistent but Modest: A Meta-analysis on Unimodal and Multimodal Affect Detection Accuracies from 30 Studies. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction (ICMI '12)*. ACM, New York, NY, USA, 31–38. DOI:http://dx.doi.org/10.1145/2388676.2388686

[40] Ellen Douglas-Cowie, Nick Campbell, Roddy Cowie, and Peter Roach. 2003. Emotional speech: Towards a new generation of databases. *Speech Communication* 40, 1–2 (2003), 33–60. DOI:http://dx.doi.org/https://doi.org/10.1016/S0167-6393(02)00070-5

[41] Ellen Douglas-Cowie, Laurence Devillers, Jean-Claude Martin, Roddy Cowie, Suzie Savvidou, Sarkis Abrilian, and Cate Cox. 2005. Multimodal databases of everyday emotion: facing up to complexity. In *Proceedings of INTERSPEECH*. 813–816.

[42] Ryan Eastman, Stephen G. Warren, Ryan Eastman, and Stephen G. Warren. 2013. A 39-Yr Survey of Cloud Changes from Land Stations Worldwide 1971–2009: Long-Term Trends, Relation to Aerosols, and Expansion of the Tropical Belt. *Journal of Climate* 26, 4 (Feb 2013), 1286–1303. DOI:http://dx.doi.org/10.1175/JCLI-D-12-00280.1

[43] Paul Ekman and Wallace V Friesen. 1971. Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology* 17 (1971), 124–129.

[44] Paul Ekman and Wallace V Friesen. 1986. A new pancultural facial expression of emotion. *Motivation and Emotion* 10 (1986), 159–168.

[45] Hillary Anger Elfenbein and Nalini Ambady. 2002. On the Universality and Cultural Specificity of Emotion Recognition: A Meta-Analysis. *Psychological Bulletin* 128, 2 (2002), 203–235.

[46] Birgit Endrass, Markus Haering, Gasser Akila, and Elisabeth André. 2014. Simulating Deceptive Cues of Joy in Humanoid Robots. In *Proceedings of the 14th International Conference on Intelligent Virtual Agents (IVA 2014)*, Timothy Bickmore, Stacy Marsella, and Candace Sidner (Eds.). Springer International Publishing, Cham, 174–177. DOI:http://dx.doi.org/10.1007/978-3-319-09767-1_20

[47] Jacinto Estima and Marco Painho. 2013. Flickr Geotagged and Publicly Available Photos: Preliminary Study of Its Adequacy for Helping Quality Control of Corine Land Cover. In *ICCSA 2013: Computational Science and Its Applications*. 205–220. DOI:http://dx.doi.org/10.1007/978-3-642-39649-6_15

[48] Florian Eyben, Martin Wöllmer, and Björn Schuller. 2009. OpenEAR: Introducing the Munich open-source emotion and affect recognition toolkit. In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*. 1–6. DOI:http://dx.doi.org/10.1109/ACII.2009.5349350

[49] Lisa Feldman Barrett, Batja Mesquita, and Maria Gendron. 2011. Context in Emotion Perception. *Current Directions in Psychological Science* 20 (2011), 286–290.

[50] José-Miguel Fernández-Dols and Carlos Crivelli. 2013. Emotion and Expression: Naturalistic Studies. *Emotion Review* 5, 1 (2013), 24–29.

[51] Gerald Friedland and Robin Sommer. 2010. Cybercasing the joint: On the privacy implications of geotagging. In *Proceedings of the 5th USENIX Conference on Hot Topics in Security (HotSec)*. USENIX Association, 1–8.

[52] Paul B. Garrett and Patricia Baquedano-López. 2002. Language socialization: Reproduction and continuity, transformation and change. *Annual Review of Anthropology* 31 (2002), 339–361.

[53] Augusta Gaspar and Francisco G. Esteves. 2012. Preschooler's faces in spontaneous emotional contexts—How well do they match adult facial expression prototypes? *International Journal of Behavioral Development* 36, 5 (2012), 348–357.

[54] Michael Grimm, Kristian Kroschel, and Shrikanth Narayanan. 2008. The Vera am Mittag German audio-visual emotional speech database. In *2008 IEEE International Conference on Multimedia and Expo*. 865–868. DOI:http://dx.doi.org/10.1109/ICME.2008.4607572

[55] Hatice Gunes and Hayley Hung. 2016. Is automatic facial expression recognition of emotions coming to a dead end? The rise of the new kids on the block. *Image and Vision Computing* 55, 1 (2016), 6–8.

[56] Hatice Gunes and Björn Schuller. 2013. Categorical and dimensional affect analysis in continuous input: Current trends and future directions. *Image and Vision Computing* 31, 2 (2013), 120–136. DOI:http://dx.doi.org/10.1016/j.imavis.2012.06.016

[57] Hatice Gunes, Björn Schuller, Maja Pantic, and Roddy Cowie. 2011. Emotion representation, analysis and synthesis in continuous space: A survey. In *Face and Gesture 2011*. 827–834. DOI:http://dx.doi.org/10.1109/FG.2011.5771357

[58] Lisa Gutzeit, Alexander Fabisch, Marc Otto, Jan Hendrik Metzen, Jonas Hansen, Frank Kirchner, and Elsa Andrea Kirchner. 2017. The BesMan Learning Platform for Automated Robot Skill Learning. *Autonomous Robots* (2017). Submitted.

[59] Lisa Gutzeit and Elsa Andrea Kirchner. 2016. Automatic Detection and Recognition of Human Movement Patterns in Manipulation Tasks. In *Proceedings of the 3rd International Conference on Physiological Computing Systems*. SCITEPRESS - Science and Technology Publications, 54–63. DOI:http://dx.doi.org/10.5220/0005946500540063

[60] Jonathan Haidt and Dacher Keltner. 1999. Culture and facial expression: Open-ended methods find more expressions and a gradient of recognition. *Cognition and Emotion* 13 (1999), 225–266.

[61] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, and Andrew Y. Ng. 2014. Deep speech: Scaling up end-to-end speech recognition. (2014). arXiv:1412.5567.

[62] Hilda Hardy, Kirk Baker, Laurence Devillers, Lori Lamel, Sophie Rosset, Cristian Ursu, and Nick Webb. 2002. Multi-layer dialogue annotation for automated multilingual customer service. In *Proceedings of the International Standards for Language Engineering Workshop*.

[63] Christine Harris and Nancy Alvarado. 2005. Facial expressions, smile types, and self-report during humor, tickle, and pain. *Cognition and Emotion* 19, 5 (2005), 655–669.

[64] Hrayr Harutyunyan, Guillaume Fidanza, and Hrant Khachatrian. Spoken language identification with deep learning. (????). https://github.com/YerevaNN/Spoken-language-identification.

[65] Koshiro Hatori. 2011. History of Origami in the East and the West before Interfusion. In *Fifth International Meeting of Origami Science, Mathematics, and Education (Origami 5)*. A K Peters/CRC Press, 3–11. DOI:http://dx.doi.org/10.1201/b10971-3 0.

[66] Thomas C Hull. 2015. Coloring connections with counting mountain-valley assignments. *Origami^6: I. Mathematics* (2015), 3.

[67] Roderick A. Hyde, Shamasundar N. Dixit, Andrew H. Weisberg, and Michael C. Rushford. 2002. Eyeglass: a very large aperture diffractive space telescope. In *SPIE 4849, Highly Innovative Space Telescope Concepts, 28*, Howard A. MacEwen (Ed.). International Society for Optics and Photonics, 28. DOI:http://dx.doi.org/10.1117/12.460420

[68] Hamid Izadinia, Bryan C. Russell, Ali Farhadi, Matthew D. Hoffman, and Aaron Hertzmann. 2015. Deep Classifiers from Image Tags in the Wild. In *Proceedings of the 2015 Workshop on Community-Organized Multimodal Mining: Opportunities for Novel Solutions (MMCommons '15)*. ACM, 13–18. DOI:http://dx.doi.org/10.1145/2814815.2814821

[69] T. Jebara and A. Pentland. 2002. Statistical imitative learning from perceptual data. In *Proceedings of the 2nd International Conference on Development and Learning (ICDL 2002)*. IEEE Comput. Soc, 191–196. DOI:http://dx.doi.org/10.1109/DEVLRN.2002.1011859

[70] Lu Jiang, Shoou-I Yu, Deyu Meng, Teruko Mitamura, and Alexander G. Hauptmann. 2015. Bridging the Ultimate Semantic Gap. In *Proceedings of the 5th ACM on International Conference on Multimedia Retrieval - ICMR '15*. ACM Press, New York, New York, USA, 27–34. DOI:http://dx.doi.org/10.1145/2671188.2749399

[71] Lu Jiang, Shoou-I Yu, Deyu Meng, Yi Yang, Teruko Mitamura, and Alexander G. Hauptmann. 2015. Fast and accurate content-based semantic search in 100M Internet videos. In *Proceedings of the 23rd ACM International Conference on Multimedia (MM '15)*. ACM, 49–58.

[72] Bálint Kádár and Mátyás Gede. 2013. Where Do Tourists Go? Visualizing and Analysing the Spatial Distribution of Geotagged Photography. *Cartographica: The International Journal for Geographic Information and Geovisualization* 48, 2 (Jun 2013), 78–88. DOI:http://dx.doi.org/10.3138/carto.48.2.1839

[73] Sebastian Kalkowski, Christian Schulze, Andreas Dengel, and Damian Borth. 2015. Real-time Analysis and Visualization of the YFCC100m Dataset. In *Proceedings of the 2015 Workshop on Community-Organized Multimodal Mining: Opportunities for Novel Solutions (MMCommons '15)*. ACM, 25–30. DOI:http://dx.doi.org/10.1145/2814815.2814820

[74] Takeo Kanade, Jeffrey F. Cohn, and Yingli Tian. 2000. Comprehensive database for facial expression analysis. In *Proceedings of the Fourth IEEE International Conference on Automatic Face and Gesture Recognition (Cat. No. PR00580)*. 46–53. DOI:http://dx.doi.org/10.1109/AFGR.2000.840611

[75] Alex Kendall, Matthew Grimes, and Roberto Cipolla. 2015. PoseNet: A Convolutional Network for Real-Time 6-DOF Camera Relocalization. In *2015 IEEE International Conference on Computer Vision (ICCV)*. IEEE, 2938–2946. DOI:http://dx.doi.org/10.1109/ICCV.2015.336

[76] Andrea Kleinsmith and Nadia Bianchi-Berthouze. 2013. Affective Body Expression Perception and Recognition: A Survey. *IEEE Transactions on Affective Computing* 4, 1 (Jan 2013), 15–33. DOI:http://dx.doi.org/10.1109/T-AFFC.2012.16

[77] Andrea Kleinsmith, Nadia Bianchi-Berthouze, and Anthony Steed. 2011. Automatic Recognition of Non-Acted Affective Postures. *IEEE Transactions on Systems, Man, and Cybernetics—Part B: Cybernetics* 41, 4 (August 2011), 1027–1038. http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=5704207

[78] Christoph Kofler, Luz Caballero, Maria Menendez, Valentina Occhialini, and Martha Larson. 2011. Near2me: an authentic and personalized social media-based recommender for travel destinations. In *Proceedings of the 3rd ACM SIGMM International Workshop on Social Media (WSM '11)*. ACM, New York, New York, USA, 47. DOI:http://dx.doi.org/10.1145/2072609.2072624

[79] Alireza Koochali, Sebastian Kalkowski, Andreas Dengel, Damian Borth, and Christian Schulze. 2016. Which Languages Do People Speak on Flickr?: A Language and Geo-Location Study of the YFCC100m Dataset. In *Proceedings of the 2016 ACM Workshop on the Multimedia COMMONS (MMCommons '16)*. ACM, 35–42.

[80] Giorgos Kordopatis-Zilos, Symeon Papadopoulos, and Yiannis Kompatsiaris. 2016. In-Depth Exploration of Geotagging Performance Using Sampling Strategies on YFCC100M. In *Proceedings of the 2016 ACM Workshop on the Multimedia COMMONS (MMCommons '16)*. ACM Press, New York, New York, USA, 3–10. DOI:http://dx.doi.org/10.1145/2983554.2983558

[81] Mario Michael Krell, Sirko Straube, Anett Seeland, Hendrik Wöhrle, Johannes Teiwes, Jan Hendrik Metzen, Elsa Andrea Kirchner, and Frank Kirchner. 2013. pySPACE—a signal processing and classification environment in Python. *Frontiers in Neuroinformatics* 7, 40 (Dec 2013), 1–11. DOI:http://dx.doi.org/10.3389/fninf.2013.00040

[82] Kaori Kuribayashi, Koichi Tsuchiya, Zhong You, Dacian Tomus, Minoru Umemoto, Takahiro Ito, and Masahiro Sasaki. 2006. Self-deployable origami stent grafts as a biomedical application of Ni-rich TiNi shape memory alloy foil. *Materials Science and Engineering: A* 419, 1–2 (Mar 2006), 131–137. DOI:http://dx.doi.org/10.1016/j.msea.2005.12.016

[83] Paul Lamere, Philip Kwok, Evandro Gouvea, Bhiksha Raj, Rita Singh, William Walker, Manfred Warmuth, and Peter Wolf. 2003. The CMU SPHINX-4 speech recognition system. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2003), Hong Kong, April 6–10, 2003*, Vol. 1. 2–5.

[84] Robert J Lang. 2008. From flapping birds to space telescopes: The modern science of origami. In *6th Symposium on Non-Photorealistic Animation and Rendering (NPAR)*. 7.

[85] Chi-Chun Lee, Carlos Busso, Sungbok Lee, and Shrikanth Narayanan. 2009. Modeling mutual influence of interlocutor emotion states in dyadic spoken interactions. In *Proceedings of INTERSPEECH*. 1983–1986.

[86] Iulia Lefter, Leon J. M. Rothkrantz, Pascal Wiggers, and David A. van Leeuwen. 2010. Emotion recognition from speech by combining databases and fusion of classifiers. In *Proceedings of the 13th International Conference on Text, Speech, and Dialogue (TSD 2010), Brno, Czech Republic, September 6–10, 2010 (Lecture Notes in Computer Science)*, Petr Sojka, Aleš Horák, Ivan Kopeček, and Karel Pala (Eds.). Springer, 353–360.

[87] Chenbin Liu, Francis Tsow, Yi Zou, and Nongjian Tao. 2016. Particle Pollution Estimation Based on Image Analysis. *PLOS ONE* 11, 2 (Feb 2016), e0145955. DOI:http://dx.doi.org/10.1371/journal.pone.0145955

[88] Marco Lui and Timothy Baldwin. 2012. Langid.Py: An Off-the-shelf Language Identification Tool. In *Proceedings of the ACL 2012 System Demonstrations (ACL '12)*. Association for Computational Linguistics, Stroudsburg, PA, USA, 25–30. http://dl.acm.org/citation.cfm?id=2390470.2390475

[89] Mathias Lux and Glenn Macstravic. 2014. The LIRE Request Handler: A Solr Plug-In for Large Scale Content Based Image Retrieval. In *MultiMedia Modeling: 20th Anniversary International Conference, MMM 2014, Dublin, Ireland, January 6-10, 2014, Proceedings, Part II*, Cathal Gurrin, Frank Hopfgartner, Wolfgang Hurst, Håvard Johansen, Hyowon Lee, and Noel O'Connor (Eds.). Springer International Publishing, Cham, 374–377. DOI:http://dx.doi.org/10.1007/978-3-319-04117-9_39

[90] Brian MacWhinney. 2000. *The CHILDES Project: Tools for Analyzing Talk* (3rd ed.). Lawrence Erlbaum Associates. http://childes.talkbank.org/.

[91] A Mammucari, C Caltagirone, P Ekman, W Friesen, G Gainotti, L Pizzamiglio, and P Zoccolotti. 1988. Spontaneous Facial Expression of Emotions in Brain-damaged Patients. *Cortex* 24 (1988), 521–533.

[92] Soroosh Mariooryad, Reza Lotfian, and Carlos Busso. 2014. Building a naturalistic emotional speech corpus by retrieving expressive behaviors from existing speech corpora. In *Proceedings of INTERSPEECH*. 238–242.

[93] Samuel Mascarenhas, Nick Degens, Ana Paiva, Rui Prada, Gert Jan Hofstede, Adrie Beulens, and Ruth Aylett. 2016. Modeling culture in intelligent virtual agents. *Autonomous Agents and Multi-Agent Systems* 30, 5 (2016), 931–962. DOI:http://dx.doi.org/10.1007/s10458-015-9312-6

[94] James Matisoff. 1986. Hearts and minds in Southeast Asian languages and English. *Cahiers de Linguistique Asie Orientale* 15, 1 (1986), 5–57.

[95] David Matsumoto and Bob Willingham. 2009. Spontaneous facial expressions of emotion of congenitally and noncongenitally blind individuals. *Journal of Personal and Social Psychology* 96, 1 (2009), 1–10. DOI:http://dx.doi.org/10.1037/a0014037

[96] Gary McKeown, Michel Valstar, Roddy Cowie, Maja Pantic, and Marc Schroder. 2012. The SEMAINE Database: Annotated Multimodal Records of Emotionally Colored Conversations Between a Person and a Limited Agent. *IEEE Transactions on Affective Computing* 3, 1 (Jan 2012), 5–17. DOI:http://dx.doi.org/10.1109/T-AFFC.2011.20

[97] Angeliki Metallinou, Martin Wollmer, Athanasios Katsamanis, Florian Eyben, Bjorn Schuller, and Shrikanth Narayanan. 2012. Context-Sensitive Learning for Enhanced Audiovisual Emotion Classification. *IEEE Transactions on Affective Computing* 3, 2 (April 2012), 184–198. DOI:http://dx.doi.org/10.1109/T-AFFC.2011.40

[98] Jan Hendrik Metzen, Alexander Fabisch, Lisa Senger, José de Gea Fernández, and Elsa Andrea Kirchner. 2013. Towards Learning of Generic Skills for Robotic Manipulation. *KI - Künstliche Intelligenz* 28, 1 (Dec 2013), 15–20. DOI:http://dx.doi.org/10.1007/s13218-013-0280-1

[99] Robert W. Mitchell. 2001. Americans' Talk to Dogs: Similarities and Differences With Talk to Infants. *Research on Language & Social Interaction* 34, 2 (Apr 2001), 183–210. DOI:http://dx.doi.org/10.1207/S15327973RLSI34-2_2

[100] Koryo Miura. 1994. Map fold a la Miura style, its physical characteristics and application to the space science. *Research of Pattern Formation* (1994), 77–90.

[101] Emily Mower, Angeliki Metallinou, Chi-Chun Lee, Abe Kazemzadeh, Carlos Busso, Sungbok Lee, and Shrikanth Narayanan. 2009. Interpreting ambiguous emotional expressions. In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*. 1–8. DOI:http://dx.doi.org/10.1109/ACII.2009.5349500

[102] Venkatesh N. Murthy, Subhransu Maji, and R. Manmatha. 2015. Automatic Image Annotation Using Deep Learning Representations. In *Proceedings of the 5th ACM International Conference on Multimedia Retrieval (ICMR '15)*. ACM, New York, NY, USA, 603–606. DOI:http://dx.doi.org/10.1145/2671188.2749391

[103] Pamela J Naab and James A Russell. 2007. Judgements of Emotion From Spontaneous Facial Expressions of New Guineans. *Emotion* 7, 4 (2007), 736–744.

[104] Elinor Ochs and Bambi B. Schieffelin. 1984. Language acquisition and socialization: Three developmental stories and their implications. In *Culture Theory: Essays on Mind, Self, and Emotion*, Richard A. Shweder and Robert A. LeVine (Eds.). Cambridge University Press, Cambridge, UK, 276–320.

[105] Elinor Ochs and Bambi B. Schieffelin. 2012. The theory of language socialization. In *The Handbook of Language Socialization*, Alessandro Duranti, Elinor Ochs, and Bambi B. Schieffelin (Eds.). Wiley-Blackwell, Malden, MA, 1–21.

[106] Şeyda Özçalışkan and Susan Goldin-Meadow. 2005. Gesture is at the cutting edge of early language development. *Cognition* 96, 3 (2005), B101–B113. DOI:http://dx.doi.org/https://doi.org/10.1016/j.cognition.2005.01.001

[107] Maja Pantic and Leon J.M. Rothkrantz. 2000. Automatic analysis of facial expressions: The state of the art. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 12 (Dec 2000), 1424–1445.

[108] Maja Pantic and Leon J.M. Rothkrantz. 2004. Case-based reasoning for user-profiled recognition of emotions from face images. In *2004 IEEE International Conference on Multimedia and Expo (ICME)*, Vol. 1. 391–394. DOI:http://dx.doi.org/10.1109/ICME.2004.1394211

[109] Maja Pantic, Michel Valstar, Ron Rademaker, and Ludo Maat. 2005. Web-based database for facial expression analysis. In *2005 IEEE International Conference on Multimedia and Expo*. DOI:http://dx.doi.org/10.1109/ICME.2005.1521424

[110] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (Feb 2011), 2825–2830. http://dl.acm.org/citation.cfm?id=1953048.2078195

[111] Adrian Popescu, Eleftherios Spyromitros-Xioufis, Symeon Papadopoulos, Hervé Le Borgne, and Ioannis Kompatsiaris. 2015. Toward an Automatic Evaluation of Retrieval Performance with Large Scale Image Collections. In *Proceedings of the 2015 Workshop on Community-Organized Multimodal Mining: Opportunities for Novel Solutions (MMCommons '15)*. ACM, 7–12. DOI:http://dx.doi.org/10.1145/2814815.2814819

[112] Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, Jan Silovsky, Georg Stemmer, and Karel Vesely. 2011. The Kaldi speech recognition toolkit. In *Proceedings of the IEEE 2011 Workshop on Automatic Speech Recognition and Understanding*. IEEE Signal Processing Society. IEEE Catalog No.: CFP11SRW-USB.

[113] Fabien Ringeval, Andreas Sonderegger, Juergen Sauer, and Denis Lalanne. 2013. Introducing the RECOLA multimodal corpus of remote collaborative and affective interactions. In *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*. 1–8. DOI:http://dx.doi.org/10.1109/FG.2013.6553805

[114] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2014. ImageNet Large Scale Visual Recognition Challenge. (Sep 2014). arXiv:1409.0575 http://arxiv.org/abs/1409.0575

[115] James A Russell. 1994. Is there universal recognition of emotion from facial expression? *Psychological Bulletin* 115, 1 (1994), 102–141.

[116] James A Russell, Jo-Anne Bachorowski, and José-Miguel Fernández-Dols. 2003. Facial and vocal expressions of emotion. *Annual Review of Psychology* 54 (2003), 329–349.

[117] Evangelos Sariyanidi, Hatice Gunes, and Andrea Cavallaro. 2015. Automatic Analysis of Facial Affect: A Survey of Registration, Representation, and Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 6 (June 2015), 1113–1133.

[118] Evangelos Sariyanidi, Hatice Gunes, Muhittin Gökmen, and Andrea Cavallaro. 2013. Local Zernike

Moment Representation for Facial Affect Recognition. In *Proceedings of the British Machine Vision Conference (BMVC)*. 108.1–108.13.

[119] Bambi B. Schieffelin and Elinor Ochs. 1986. Language socialization. *Annual Review of Anthropology* 15 (1986), 163–191.

[120] Björn Schuller, Dino Seppi, Anton Batliner, Andreas Maier, and Stefan Steidl. 2007. Towards More Reality in the Recognition of Emotional Speech. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP '07)*, Vol. 4. IV.941–944. DOI:http://dx.doi.org/10.1109/ICASSP.2007.367226

[121] Björn Schuller, Bogdan Vlasenko, Florian Eyben, Martin Wollmer, Andre Stuhlsatz, Andreas Wendemuth, and Gerhard Rigoll. 2010. Cross-Corpus Acoustic Emotion Recognition: Variances and Strategies. *IEEE Transactions on Affective Computing* 1, 2 (July 2010), 119–131. DOI:http://dx.doi.org/10.1109/T-AFFC.2010.8

[122] Lisa Senger, Martin Schroer, Jan Hendrik Metzen, and Elsa Andrea Kirchner. 2014. Velocity-Based Multiple Change-Point Inference for Unsupervised Segmentation of Human Movement Behavior. In *2014 22nd International Conference on Pattern Recognition*. IEEE, 4564–4569. DOI:http://dx.doi.org/10.1109/ICPR.2014.781

[123] Pierre Sermanet and Yann LeCun. 2011. Traffic sign recognition with multi-scale Convolutional Networks. In *The 2011 International Joint Conference on Neural Networks*. 2809–2813. DOI:http://dx.doi.org/10.1109/IJCNN.2011.6033589

[124] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations (ICLR 2015)*. arXiv:1409.1556 http://arxiv.org/abs/1409.1556

[125] Sirko Straube and Mario Michael Krell. 2014. How to evaluate an agent's behaviour to infrequent events? – Reliable performance estimation insensitive to class distribution. *Frontiers in Computational Neuroscience* 8, 43 (Jan 2014), 1–6. DOI:http://dx.doi.org/10.3389/fncom.2014.00043

[126] Patricia L. Sunderland and Rita M. Denny. 2016. *Doing Anthropology in Consumer Research*. Routledge.

[127] Bart Thomee, Benjamin Elizalde, David A. Shamma, Karl Ni, Gerald Friedland, Douglas Poland, Damian Borth, and Li-Jia Li. 2016. YFCC100M: The New Data in Multimedia Research. *Communications of the ACM* 59, 2 (Jan 2016), 64–73. DOI:http://dx.doi.org/10.1145/2812802

[128] Michel F. Valstar and Maja Pantic. 2012. Fully Automatic Recognition of the Temporal Phases of Facial Actions. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 42, 1 (Feb 2012), 28–43. DOI:http://dx.doi.org/10.1109/TSMCB.2011.2163710

[129] Kenneth Van-Bik. 1998. Lai Psycho-collocation. *Linguistics of the Tibeto-Burman Area* 21, 1 (1998), 201–233.

[130] Various. 2017. Origami Database: http://www.oriwiki.com/. (2017).

[131] Laurence Vidrascu and Laurence Devillers. 2008. Anger detection performances based on prosodic and acoustic cues in several corpora. In *Proceedings of the Workshop on Corpora for Research on Emotion and Affect at the 2008 International Conference on Language Resources and Evaluation (LREC 2008)*. 13–16.

[132] Thurid Vogt and Elisabeth André. 2005. Comparing Feature Sets for Acted and Spontaneous Speech in View of Automatic Emotion Recognition. In *2005 IEEE International Conference on Multimedia and Expo (ICME '05)*. 474–477. DOI:http://dx.doi.org/10.1109/ICME.2005.1521463

[133] Thurid Vogt, Elisabeth André, and Nikolaus Bee. 2008. EmoVoice: A framework for online recognition of emotions from voice. In *Proceedings of Workshop on Perception and Interactive Technologies for Speech-Based Systems*.

[134] Thurid Vogt, Elisabeth André, and Johannes Wagner. 2008. Automatic Recognition of Emotions from Speech: A Review of the Literature and Recommendations for Practical Realisation. In *Affect and Emotion in Human-Computer Interaction: From Theory to Applications*, Christian Peter and Russell Beale (Eds.). Springer, Berlin/Heidelberg, 75–91. DOI:http://dx.doi.org/10.1007/978-3-540-85099-1_7

[135] Jingya Wang, Mohammed Korayem, Saul Blanco, and David J. Crandall. 2016. Tracking Natural Events through Social Media and Computer Vision. In *Proceedings of the 2016 ACM Conference on Multimedia - MM '16*. ACM Press, New York, New York, USA, 1097–1101. DOI:http://dx.doi.org/10.1145/2964284.2984067

[136] Yilin Wang, Jiliang Tang, Jundong Li, Baoxin Li, Yali Wan, Clayton Mellina, Neil O Hare, and Yi Chang.

2017. Understanding and Discovering Deliberate Self-harm Content in Social Media. In *International World Wide Web Conference (WWW)*.

[137] Linda R. Watson, Elizabeth R. Crais, Grace T. Baranek, Jessica R. Dykstra, and Kaitlyn P. Wilson. 2013. Communicative gesture use in infants with and without autism: A retrospective home video study. *American Journal of Speech-Language Pathology* 22, 1 (2013), 25–39.

[138] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. 2016. Convolutional Pose Machines. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. arXiv:1602.00134 http://arxiv.org/abs/1602.00134

[139] John Wiseman and Ivan Yu. Bondarenko. Python interface to the WebRTC Voice Activity Detector. (????). https://github.com/wiseman/py-webrtcvad.

[140] Min Xia, Weitao Lu, Jun Yang, Ying Ma, Wen Yao, and Zichen Zheng. 2015. A hybrid method based on extreme learning machine and k-nearest neighbor for cloud classification of ground-based visible cloud image. *Neurocomputing* 160 (Jul 2015), 238–249. DOI:http://dx.doi.org/10.1016/j.neucom.2015.02.022

[141] Shicheng Xu, Susanne Burger, Alexander Hauptmann, Huan Li, Xiaojun Chang, Shoou-I Yu, Xingzhong Du, Xuanchong Li, Lu Jiang, Zexi Mao, and Zhenzhong Lan. 2015. Incremental Multimodal Query Construction for Video Search. In *Proceedings of the 5th ACM on International Conference on Multimedia Retrieval - ICMR '15*. ACM Press, New York, New York, USA, 675–678. DOI:http://dx.doi.org/10.1145/2671188.2749413

[142] Yezhou Yang, Yi Li, Cornelia Fermüller, and Yiannis Aloimonos. 2015. Robot Learning Manipulation Action Plans by 'Watching' Unconstrained Videos from the World Wide Web. In *Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*.

[143] Hashim Yasin, Umar Iqbal, Bjorn Kruger, Andreas Weber, and Juergen Gall. 2016. A Dual-Source Approach for 3D Pose Estimation from a Single Image. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 4948–4956. DOI:http://dx.doi.org/10.1109/CVPR.2016.535 arXiv:1509.06720

[144] Zhihong Zeng, Maja Pantic, Glenn I. Roisman, and Thomas S. Huang. 2009. A Survey of Affect Recognition Methods: Audio, Visual, and Spontaneous Expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31, 1 (Jan 2009), 39–58. DOI:http://dx.doi.org/10.1109/TPAMI.2008.52

[145] Chao Zhang, Junchi Yan, Changsheng Li, Xiaoguang Rui, Liang Liu, and Rongfang Bie. 2016. On Estimating Air Pollution from Photos Using Convolutional Neural Network. In *Proceedings of the 2016 ACM Conference on Multimedia (MM'16)*. ACM, 297–301.

[146] Haipeng Zhang, Mohammed Korayem, David J. Crandall, and Gretchen LeBuhn. 2012. Mining photo-sharing websites to study ecological phenomena. In *Proceedings of the 21st International Conference on World Wide Web (WWW '12)*. ACM Press, New York, New York, USA, 749. DOI:http://dx.doi.org/10.1145/2187836.2187938

[147] Xiaowei Zhou, Menglong Zhu, Spyridon Leonardos, Konstantinos G. Derpanis, and Kostas Daniilidis. 2016. Sparseness Meets Deepness: 3D Human Pose Estimation From Monocular Video. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.