

# Topics Seminar Media Informatics

Updated: March 9, 2020

Jörg Cassens

Seminar Medieninformatik

Winter Term 2019/2020



## Topics

A list of texts not assigned yet can be found on page 19.

### 1. Natural Language: Understanding and Generation

#### 1.1. Predicting Quality and Interestingness of Texts

□ **Topic not assigned**

- ▷ Paper: Debasis Ganguly, Johannes Leveling, and Gareth J.F. Jones. Automatic prediction of text aesthetics and interestingness. In Proceedings of the International Conference on Computational Linguistics. 2014.

[Download paper](#)

*Abstract: This paper investigates the problem of automated text aesthetics prediction. The availability of user generated content and ratings, e.g. Flickr, has induced research in aesthetics prediction for non-text domains, particularly for photographic images. This problem, however, has yet not been explored for the text domain. Due to the very subjective nature of text aesthetics, it is difficult to compile human annotated data by methods such as crowd sourcing with a fair degree of inter-annotator agreement. The availability of the Kindle “popular highlights” data has motivated us to compile a dataset comprised of human annotated aesthetically pleasing and interesting text passages. We then undertake a supervised classification approach to predict text aesthetics by constructing real-valued feature vectors from each text passage. In particular, the features that we use for this classification task are word length, repetitions, polarity, part-of-speech, semantic distances; and topic generality and diversity. A traditional binary classification approach is not effective in this case because non-highlighted passages surrounding the highlighted ones do not necessarily represent the other extreme of unpleasant quality text. Due to the absence of real negative class samples, we employ the MC algorithm, in which training can be initiated with instances only from the positive class. On each successive iteration the algorithm selects new strong negative samples from the unlabeled class and retrains itself. The results show that the mapping convergence (MC) algorithm with a Gaussian and a linear kernel used for the mapping and convergence phases, respectively, yields the best results, achieving satisfactory accuracy, precision and recall values of about 74%, 42% and 54% respectively.*

- ▷ Paper: Tong Wang, Ping Chen, Boyang Li: Predicting the Quality of Short Narratives from Social Media. Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence. Main track. Pages 3859-3865.

[Download paper](#)

*Abstract: An important and difficult challenge in building computational models for narratives is the automatic evaluation of narrative quality. Quality evaluation connects narrative understanding and generation as generation systems need to evaluate their own products. To circumvent difficulties in acquiring annotations, we employ upvotes in social media as an approximate measure for story quality. We collected 54,484 answers from a crowd-powered question-and-answer website, Quora, and then used active learning to build a classifier that labeled 28,320 answers as stories. To predict the number of upvotes without the use of social network features, we create neural networks that model textual regions and the interdependence among regions, which serve as strong benchmarks for future research. To our best knowledge, this is the first large-scale study for automatic evaluation of narrative quality.*

## 1.2. **Headline Generation**

☒ **Presented by: Lea Bücher**

- ▷ Paper: Enrique Alfonseca, Daniele Pighin, and Guillermo Garrido. HEADY: news headline abstraction through event pattern clustering. In ACL 2013, pages 1243–1253, 2013.

[Download paper](#)

Abstract: *This paper presents HEADY: a novel, abstractive approach for headline generation from news collections. From a web-scale corpus of English news, we mine syntactic patterns that a Noisy-OR model generalizes into event descriptions. At inference time, we query the model with the patterns observed in an unseen news collection, identify the event that better captures the gist of the collection and retrieve the most appropriate pattern to generate a headline. HEADY improves over a state-of-the-art open-domain title abstraction method, bridging half of the gap that separates it from extractive methods using human-generated titles in manual evaluations, and performs comparably to human-generated headlines as evaluated with ROUGE.*

- ▷ Paper: Jiwei Tan, Xiaojun Wan, Jianguo Xiao: From Neural Sentence Summarization to Headline Generation: A Coarse-to-Fine Approach. Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence. Main track. Pages 4109-4115.

[Download paper](#)

Abstract: *Headline generation is a task of abstractive text summarization, and previously suffers from the immaturity of natural language generation techniques. Recent success of neural sentence summarization models shows the capacity of generating informative, fluent headlines conditioned on selected recapitulative sentences. In this paper, we investigate the extension of sentence summarization models to the document headline generation task. The challenge is that extending the sentence summarization model to consider more document information will mostly confuse the model and hurt the performance. In this paper, we propose a coarse-to-fine approach, which first identifies the important sentences of a document using document summarization techniques, and then exploits a multi-sentence summarization model with hierarchical attention to leverage the important sentences for headline generation. Experimental results on a large real dataset demonstrate the proposed approach significantly improves the performance of neural sentence summarization models on the headline generation task.*

## 1.3. **Time Aware Knowledge Extraction for Summarization**

☐ **Topic not assigned**

- ▷ Paper: Carmen De Maio, Giuseppe Fenza, Vincenzo Loia, Mimmo Parente, Time Aware Knowledge Extraction for microblog summarization on Twitter, In Information Fusion, Volume 28, 2016, Pages 60-74, ISSN 1566-2535

[Download paper](#)

Abstract: *Microblogging services like Twitter and Facebook collect millions of user generated content every moment about trending news, occurring events, and so on. Nevertheless, it is really a nightmare to find information of interest through the huge amount of available posts that are often noisy and redundant. In the era of Big Data, social media analytics services have caught increasing attention from both research and industry. Specifically, the dynamic context of microblogging requires to manage not only meaning of information but also the evolution of knowledge over the timeline. This work defines Time Aware Knowledge Extraction (briefly TAKE) methodology that relies on temporal extension of Fuzzy Formal Concept Analysis. In particular, a microblog summarization algorithm has been defined filtering the concepts organized by TAKE in a time-dependent hierarchy. The algorithm addresses topic-based summarization on Twitter. Besides considering the timing of the concepts, another distinguishing feature of the proposed microblog summarization framework is the possibility to have more or less detailed summary, according to the user's needs, with good levels of quality and completeness as highlighted in the experimental results.*

## 1.4. **Annotator Rationales for Relevance Judgment**

☐ **Topic not assigned**

- ▷ Paper: Tyler McDonnell, Matthew Lease, Mucahid Kutlu, and Tamer Elsayed. Why Is That Relevant? Collecting Annotator Rationales for Relevance Judgments. In Proc. of the 4th AAI Conference on Human Computation and Crowdsourcing (HCOMP), pages 139–148, 2016.

[Download paper](#)

Abstract: *When collecting subjective human ratings of items, it can be difficult to measure and enforce data quality due to task subjectivity and lack of insight into how judges arrive at each rating decision. To address this, we propose requiring judges to provide a specific type of rationale underlying each rating decision. We evaluate this approach in the domain of Information Retrieval, where human judges rate the relevance of Webpages to search queries. Cost-benefit analysis over 10,000 judgments collected on Mechanical Turk suggests a win-win: experienced crowd workers provide rationales with almost no increase in task completion time while providing a multitude of further benefits, including more reliable*

judgments and greater transparency for evaluating both human raters and their judgments. Further benefits include reduced need for expert gold, the opportunity for dual-supervision from ratings and rationales, and added value from the rationales themselves.

- ▷ Paper: Tyler McDonnell, Mucahid Kutlu, Tamer Elsayed, Matthew Lease: The Many Benefits of Annotator Rationales for Relevance Judgments. Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence Best Sister Conferences. Pages 4909-4913

[Download paper](#)

Abstract: *When collecting subjective human ratings of items, it can be difficult to measure and enforce data quality due to task subjectivity and lack of insight into how judges arrive at each rating decision. To address this, we propose requiring judges to provide a specific type of rationale underlying each rating decision. We evaluate this approach in the domain of Information Retrieval, where human judges rate the relevance of Webpages. Cost-benefit analysis over 10,000 judgments collected on Mechanical Turk suggests a win-win: experienced crowd workers provide rationales with no increase in task completion time while providing further benefits, including more reliable judgments and greater transparency.*

## 2. Affective Computing

### 2.1. Survey Affective Computing

□ **Topic not assigned**

- ▷ Paper: Poria, S., Cambria, E., Bajpai, R., & Hussain, A. (2017). A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion*, 37, 98-125.

[Download paper](#)

Abstract: *Affective computing is an emerging interdisciplinary research field bringing together researchers and practitioners from various fields, ranging from artificial intelligence, natural language processing, to cognitive and social sciences. With the proliferation of videos posted online (e.g., on YouTube, Facebook, Twitter) for product reviews, movie reviews, political views, and more, affective computing research has increasingly evolved from conventional unimodal analysis to more complex forms of multimodal analysis. This is the primary motivation behind our first of its kind, comprehensive literature review of the diverse field of affective computing. Furthermore, existing literature surveys lack a detailed discussion of state of the art in multimodal affect analysis frameworks, which this review aims to address. Multimodality is defined by the presence of more than one modality or channel, e.g., visual, audio, text, gestures, and eye gage. In this paper, we focus mainly on the use of audio, visual and text information for multimodal affect analysis, since around 90% of the relevant literature appears to cover these three modalities. Following an overview of different techniques for unimodal affect analysis, we outline existing methods for fusing information from different modalities. As part of this review, we carry out an extensive study of different categories of state-of-the-art fusion techniques, followed by a critical analysis of potential performance improvements with multimodal analysis compared to unimodal analysis. A comprehensive overview of these two complementary fields aims to form the building blocks for readers, to better understand this challenging and exciting research field.*

### 2.2. Survey emotional body gesture recognition

□ **Topic not assigned**

- ▷ Paper: Noroozi, Fatemeh, Ciprian Adrian Corneanu, Dorota Kamińska, Tomasz Sapiński, Sergio Escalera, and Gholamreza Anbarjafari. "Survey on emotional body gesture recognition." arXiv preprint arXiv:1801.07481 (2018).

[Download paper](#)

Abstract: *Automatic emotion recognition has become a trending research topic in the past decade. While works based on facial expressions or speech abound recognizing affect from body gestures remains a less explored topic. We present a new comprehensive survey hoping to boost research in the field. We first introduce emotional body gestures as a component of what is commonly known as "body language" and comment general aspects as gender differences and culture dependence. We then define a complete framework for automatic emotional body gesture recognition. We introduce person detection and comment static and dynamic body pose estimation methods both in RGB and 3D. We then comment the recent literature related to representation learning and emotion recognition from images of emotionally expressive gestures. We also discuss multi-modal approaches that combine speech or face with body gestures for improved emotion recognition. While pre-processing methodologies (e.g. human detection and pose estimation) are nowadays mature technologies fully developed for robust large scale analysis, we show that for emotion recognition the quantity of labelled data is scarce, there is no agreement on clearly defined output spaces and the representations are shallow and largely based on naive geometrical representations.*

### 2.3. Affective Loop

□ **Topic not assigned**

- ▷ Paper: Höök K. (2008) Affective Loop Experiences – What Are They?. In: Oinas-Kukkonen H., Hasle P., Harjumaa M., Segerstahl K., Øhrstrøm P. (eds) Persuasive Technology. PERSUASIVE 2008. Lecture Notes in Computer Science, vol 5033. Springer, Berlin, Heidelberg

[Download paper](#)

Abstract: *A research agenda for bodily persuasion through a design approach we name affective loops is outlined. Affective loop experiences draw upon physical, emotional interactions between user and system.*

### 2.4. Engagement during learning

□ **Topic not assigned**

- ▷ Paper: D'Mello, Sidney, Ed Dieterle, and Angela Duckworth. Advanced, analytic, automated (AAA) measurement of engagement during learning. Educational psychologist 52, no. 2 (2017): 104-123.

[Download paper](#)

Abstract: *It is generally acknowledged that engagement plays a critical role in learning. Unfortunately, the study of engagement has been stymied by a lack of valid and efficient measures. We introduce the advanced, analytic, and automated (AAA) approach to measure engagement at fine-grained temporal resolutions. The AAA measurement approach is grounded in embodied theories of cognition and affect, which advocate a close coupling between thought and action. It uses machine-learned computational models to automatically infer mental states associated with engagement (e.g., interest, flow) from machine-readable behavioral and physiological signals (e.g., facial expressions, eye tracking, click-stream data) and from aspects of the environmental context. We present 15 case studies that illustrate the potential of the AAA approach for measuring engagement in digital learning environments. We discuss strengths and weaknesses of the AAA approach, concluding that it has significant promise to catalyze engagement research.*

### 2.5. Emotion Recognition and Adaptive User Interaction

☒ **Presented by: Pinto Dominguez Enrique**

- ▷ Paper: Meudt, Sascha, Miriam Schmidt-Wack, Frank Honold, Felix Schüssel, Michael Weber, Friedhelm Schwenker, and Günther Palm. Going further in affective computing: how emotion recognition can improve adaptive user interaction. In Toward Robotic Socially Believable Behaving Systems-Volume I, pp. 73-103. Springer, Cham, 2016.

[Download paper](#)

Abstract: *This article joins the fields of emotion recognition and human computer interaction. While much work has been done on recognizing emotions, they are hardly used to improve a user's interaction with a system. Although the fields of affective computing and especially serious games already make use of detected emotions, they tend to provide application and user specific adaptations only on the task level. We present an approach of utilizing recognized emotions to improve the interaction itself, independent of the underlying application at hand. Examining the state of the art in emotion recognition research and based on the architecture of Companion-System, a generic approach for determining the main cause of an emotion within the history of interactions is presented, allowing a specific reaction and adaptation. Using such an approach could lead to systems that use emotions to improve not only the outcome of a task but the interaction itself in order to be truly individual and empathic.*

### 2.6. Psychophysiology in Games

☒ **Presented by: Vildan Uzunçakmak**

- ▷ Paper: Yannakakis G.N., Martinez H.P., Garbarino M. (2016) Psychophysiology in Games. In: Karpouzis K., Yannakakis G. (eds) Emotion in Games. Socio-Affective Computing, vol 4. Springer, Cham

[Download paper](#)

Abstract: *Psychophysiology is the study of the relationship between psychology and its physiological manifestations. That relationship is of particular importance for both game design and ultimately game-playing. Players' psychophysiology offers a gateway towards a better understanding of playing behavior and experience. That knowledge can, in turn, be beneficial for the player as it allows designers to make better games for them; either explicitly by altering the game during play or implicitly during the game design process. This chapter argues for the importance of physiology for the investigation of player affect in games, reviews the current state of the art in sensor technology and outlines the key phases for the application of psychophysiology in games.*

## 2.7. Affective Level Design

☒ **Presented by: Till Strzyso**

- ▷ Paper: Balducci, F., Grana, C. and Cucchiara, R., 2017. Affective level design for a role-playing videogame evaluated by a brain-computer interface and machine learning methods. *The Visual Computer*, 33(4), pp.413-427.

[Download paper](#)

*Abstract: Game science has become a research field, which attracts industry attention due to a worldwide rich sell-market. To understand the player experience, concepts like flow or boredom mental states require formalization and empirical investigation, taking advantage of the objective data that psychophysiological methods like electroencephalography (EEG) can provide. This work studies the affective ludology and shows two different game levels for Neverwinter Nights 2 developed with the aim to manipulate emotions; two sets of affective design guidelines are presented, with a rigorous formalization that considers the characteristics of role-playing genre and its specific gameplay. An empirical investigation with a brain-computer interface headset has been conducted: by extracting numerical data features, machine learning techniques classify the different activities of the gaming sessions (task and events) to verify if their design differentiation coincides with the affective one. The observed results, also supported by subjective questionnaires data, confirm the goodness of the proposed guidelines, suggesting that this evaluation methodology could be extended to other evaluation tasks.*

## 2.8. Modeling Motivation in a Social Network Game

☒ **Presented by: Mark Erhardt**

- ▷ Paper: Birk, Max V., Dereck Toker, Regan L. Mandryk, and Cristina Conati. Modeling motivation in a social network game using player-centric traits and personality traits. In *International Conference on User Modeling, Adaptation, and Personalization*, pp. 18-30. Springer, Cham, 2015.

[Download paper](#)

*Abstract: People are drawn to play different types of videogames and find enjoyment in a range of gameplay experiences. Envisaging a representative game player or persona allows game designers to personalize game content; however, there are many ways to characterize players and little guidance on which approaches best model player behavior and preference. To provide knowledge about how player characteristics contribute to game experience, we investigate how personality traits as well as player styles from the BrianHex model moderate the prediction of player motivation with a social network game. Our results show that several player characteristics impact motivation, expressed in terms of enjoyment and effort. We also show that player enjoyment and effort, as predicted by our models, impact players' in-game behaviors, illustrating both the predictive power and practical utility of our models for guiding user adaptation.*

## 2.9. Identify Affective Expressions in Educational Games

□ **Topic not assigned**

- ▷ Paper: Amershi, Saleema, Cristina Conati, and Heather Maclaren. Using feature selection and unsupervised clustering to identify affective expressions in educational games. (2016).

[Download paper](#)

*Abstract: Educational games can induce a wide range of emotions, and so recognizing specific emotions may be valuable for an intelligent system that aims to adapt to varying student needs so as to improve learning. The long-term goal of this work is to understand how user affect impacts overall learning in an educational game. The main contribution of this paper is an investigation into the use of an unsupervised machine learning technique to help recognize meaningful patterns in biometric affective data. Results show that this method can identify interesting and sensible student reactions to different game events*

## 2.10. Artistic Brain-Computer Interfaces

☒ **Presented by: Aleksei Reveriuk**

- ▷ Paper: Andujar, Marvin, Chris S. Crawford, Anton Nijholt, France Jackson, and Juan E. Gilbert. Artistic brain-computer interfaces: the expression and stimulation of the user's affective state. *Brain-computer interfaces* 2, no. 2-3 (2015): 60-69.

[Download paper](#)

*Abstract: Science, technology, engineering, and mathematics (STEM) is rapidly transitioning to STEAM, which is the integration of the arts and design into the sciences. This transition is due to the need of adapting creativity in the sciences and engineering fields. This demonstrates the growing importance of the arts contribution towards the sciences in various ways. An area where art shows promise is in brain-computer interfaces. In this paper, we propose a definition for artistic brain-computer interfaces (artistic BCI) from a passive BCI perspective. We defined its composition in four fields: human-computer interaction, neurophysiology, art, and computing. This definition is based on the user's state stimulation and expression and how art could help us and the end user understand the user's affective state. We also*

discuss its applicability towards different areas and how these areas can benefit from artistic BCI.

### 3. Multimodality

#### 3.1. Multimodal Interaction

☒ **Presented by: Denise Janßen**

▷ Paper: M Turk: Multimodal interaction: A review. Pattern Recognition Letters, Volume 36, January, 2014, Pages 189-195. Elsevier

[Download paper](#)

*Abstract: People naturally interact with the world multimodally, through both parallel and sequential use of multiple perceptual modalities. Multimodal human-computer interaction has sought for decades to endow computers with similar capabilities, in order to provide more natural, powerful, and compelling interactive experiences. With the rapid advance in non-desktop computing generated by powerful mobile devices and affordable sensors in recent years, multimodal research that leverages speech, touch, vision, and gesture is on the rise. This paper provides a brief and personal review of some of the key aspects and issues in multimodal interaction, touching on the history, opportunities, and challenges of the area, especially in the area of multimodal integration. We review the question of early vs. late integration and find inspiration in recent evidence in biological sensory integration. Finally, we list challenges that lie ahead for research in multimodal human-computer interaction.*

▷ Paper: MR Morris: Web on the wall: insights from a multimodal interaction elicitation study. Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces, Pages 95-104, 2012. ACM

[Download paper](#)

*Abstract: New sensing technologies like Microsoft's Kinect provide a low-cost way to add interactivity to large display surfaces, such as TVs. In this paper, we interview 25 participants to learn about scenarios in which they would like to use a web browser on their living room TV. We then conduct an interaction-elicitation study in which users suggested speech and gesture interactions for fifteen common web browser functions. We present the most popular suggested interactions, and supplement these findings with observational analyses of common gesture and speech conventions adopted by our participants. We also reflect on the design of multimodal, multi-user interaction-elicitation studies, and introduce new metrics for interpreting user-elicitation study findings.*

#### 3.2. Gesture & Speech

☒ **Presented by: Anna Badalyan**

▷ Paper: P Wagner, Z Malisz, S Kopp: Gesture and speech in interaction: An overview. In: Journal Speech Communication, Volume 57, February, 2014. Pages 209-232. Elsevier

[Download paper](#)

*Abstract: Gestures and speech interact. They are linked in language production and perception, with their interaction contributing to felicitous communication. The multifaceted nature of these interactions has attracted considerable attention from the speech and gesture community. This article provides an overview of our current understanding of manual and head gesture form and function, of the principle functional interactions between gesture and speech aiding communication, transporting meaning and producing speech. Furthermore, we present an overview of research on temporal speech-gesture synchrony, including the special role of prosody in speech-gesture alignment. In addition, we provide a summary of tools and data available for gesture analysis, and describe speech-gesture interaction models and simulations in technical systems. This overview also serves as an introduction to a Special Issue covering a wide range of articles on these topics. We provide links to the Special Issue throughout this paper.*

#### 3.3. Multimedia Summarization

□ **Topic not assigned**

▷ Paper: F Metze, D Ding, E Younessian, A Hauptmann: Beyond audio and video retrieval: topic-oriented multimedia summarization. In: International Journal of Multimedia Information (2013) Vol. 2, pages 131-144. Springer

[Download paper](#)

*Abstract: Given the deluge of multimedia content that is becoming available over the Internet, it is increasingly important to be able to effectively examine and organize these large stores of information in ways that go beyond browsing or collaborative filtering. In this paper we review previous work on audio and video processing, and define the task of Topic-Oriented Multimedia Summarization (TOMS) using natural language generation: given a set of automatically extracted features from a video (such as visual concepts and ASR transcripts) a TOMS system will automatically generate a paragraph of natural language ("a recounting"), which summarizes the important information in a video belonging to a certain topic area, and provides explanations for why a video was matched and retrieved. We see*

this as a first step towards systems that will be able to discriminate visually similar, but semantically different videos, compare two videos and provide textual output or summarize a large number of videos at once. In this paper, we introduce our approach of solving the TOMS problem. We extract visual concept features and ASR transcription features from a given video, and develop a template-based natural language generation system to produce a textual recounting based on the extracted features. We also propose possible experimental designs for continuously evaluating and improving TOMS systems, and present results of a pilot evaluation of our initial system.

### 3.4. Multimodal Ensembles

□ **Topic not assigned**

- ▷ Paper: Maskey, S., & Hirschberg, J. (2005). Comparing lexical, acoustic/prosodic, structural and discourse features for speech summarization. In Ninth European Conference on Speech Communication and Technology.

[Download paper](#)

Abstract: We present results of an empirical study of the usefulness of different types of features in selecting extractive summaries of news broadcasts for our Broadcast News Summarization System. We evaluate lexical, prosodic, structural and discourse features as predictors of those news segments which should be included in a summary. We show that a summarization system that uses a combination of these feature sets produces the most accurate summaries, and that a combination of acoustic/prosodic and structural features are enough to build a "good" summarizer when speech transcription is not available.

- ▷ Paper: Sameer Maskey and Julia Hirschberg. 2006. Summarizing speech without text using hidden Markov models. In Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers (NAACL-Short '06). Association for Computational Linguistics, Stroudsburg, PA, USA, 89-92.

[Download paper](#)

Abstract: We present a method for summarizing speech documents without using any type of transcript/text in a Hidden Markov Model framework. The hidden variables or states in the model represent whether a sentence is to be included in a summary or not, and the acoustic/prosodic features are the observation vectors. The model predicts the optimal sequence of segments that best summarize the document. We evaluate our method by comparing the predicted summary with one generated by a human summarizer. Our results indicate that we can generate 'good' summaries even when using only acoustic/prosodic information, which points toward the possibility of text-independent summarization for spoken documents.

- ▷ Paper: Levitan, S. I., An, G., Ma, M., Levitan, R., Rosenberg, A., & Hirschberg, J. (2016). Combining Acoustic-Prosodic, Lexical, and Phonotactic Features for Automatic Deception Detection. In INTERSPEECH (pp. 2006-2010).

[Download paper](#)

Abstract: Improving methods of automatic deception detection is an important goal of many researchers from a variety of disciplines, including psychology, computational linguistics, and criminology. We present a system to automatically identify deceptive utterances using acoustic-prosodic, lexical, syntactic, and phonotactic features. We train and test our system on the Interspeech 2016 ComParE challenge corpus, and find that our combined features result in performance well above the challenge baseline on the development data. We also perform feature ranking experiments to evaluate the usefulness of each of our feature sets. Finally, we conduct a cross-corpus evaluation by training on another deception corpus and testing on the ComParE corpus.

- ▷ Paper: Mendels, G., Levitan, S. I., Lee, K. Z., & Hirschberg, J. (2017). Hybrid Acoustic-Lexical Deep Learning Approach for Deception Detection. Proc. Interspeech 2017, 1472-1476.

[Download paper](#)

Abstract: Automatic deception detection is an important problem with far-reaching implications for many disciplines. We present a series of experiments aimed at automatically detecting deception from speech. We use the Columbia X-Cultural Deception (CXD) Corpus, a large-scale corpus of within-subject deceptive and non-deceptive speech, for training and evaluating our models. We compare the use of spectral, acoustic-prosodic, and lexical feature sets, using different machine learning models. Finally, we design a single hybrid deep model with both acoustic and lexical features trained jointly that achieves state-of-the-art results on the CXD corpus.

### 3.5. Multimodal Analysis, Data Mining, Visualization

□ **Topic not assigned**

- ▷ Paper: O'Halloran, Kay L., Sabine Tan, Duc-Son Pham, John Bateman, and Andrew Vande Movere. "A digital mixed methods research design: Integrating multimodal analysis with data mining and information visualization for big data analytics." Journal of Mixed Methods Research 12, no. 1 (2018): 11-30.

### [Download paper](#)

*Abstract: This article demonstrates how a digital environment offers new opportunities for transforming qualitative data into quantitative data in order to use data mining and information visualization for mixed methods research. The digital approach to mixed methods research is illustrated by a framework which combines qualitative methods of multimodal discourse analysis with quantitative methods of data mining and information visualization in a multilevel, contextual model that will result in an integrated, theoretically well-founded, and empirically evaluated technology for analyzing large data sets of multimodal texts. The framework is applicable to situations in which critical information needs to be extracted from geotagged public data: for example, in crisis informatics, where public reports of extreme events provide valuable data sources for disaster management.*

### 3.6. Predicting Dialogue Acts for Intelligent Agents

#### □ **Topic not assigned**

- ▷ Paper: Min, Wookhee, Joseph B. Wiggins, Lydia G. Pezzullo, Alexandria K. Vail, Kristy Elizabeth Boyer, Bradford W. Mott, Megan H. Frankosky, Eric N. Wiebe, and James C. Lester. "Predicting Dialogue Acts for Intelligent Virtual Agents with Multimodal Student Interaction Data." International Educational Data Mining Society (2016).

### [Download paper](#)

*Abstract: Recent years have seen a growing interest in intelligent game-based learning environments featuring virtual agents. A key challenge posed by incorporating virtual agents in game-based learning environments is dynamically determining the dialogue moves they should make in order to best support students' problem solving. This paper presents a data-driven modeling approach that uses a Wizard-of-Oz framework to predict human wizards' dialogue acts based on a sequence of multimodal data streams of student interactions with a game-based learning environment. To effectively deal with multiple, parallel sequential data streams, this paper investigates two sequence-labeling techniques: long short-term memory networks (LSTMs) and conditional random fields. We train predictive models utilizing data corpora collected from two Wizard-of-Oz experiments in which a human wizard played the role of the virtual agent unbeknownst to the student. Empirical results suggest that LSTMs that utilize game trace logs and facial action units achieve the highest predictive accuracy. This work can inform the design of intelligent virtual agents that leverage rich multimodal student interaction data in game-based learning environments.*

## 4. Understanding & Enhancing Video

### 4.1. Visual Lecture Transcripts

#### □ **Topic not assigned**

- ▷ Paper: V Shin, F Berthouzoz, W Li, F Durand: Visual transcripts: lecture notes from blackboard-style lecture videos. In: ACM Transactions on Graphics (TOG), Volume 34 Issue 6, November 2015, pages 1-10. ACM

### [Download paper](#)

*Abstract: Blackboard-style lecture videos are popular, but learning using existing video player interfaces can be challenging. Viewers cannot consume the lecture material at their own pace, and the content is also difficult to search or skim. For these reasons, some people prefer lecture notes to videos. To address these limitations, we present Visual Transcripts, a readable representation of lecture videos that combines visual information with transcript text. To generate a Visual Transcript, we first segment the visual content of a lecture into discrete visual entities that correspond to equations, figures, or lines of text. Then, we analyze the temporal correspondence between the transcript and visuals to determine how sentences relate to visual entities. Finally, we arrange the text and visuals in a linear layout based on these relationships. We compare our result with a standard video player, and a state-of-the-art interface designed specifically for blackboard-style lecture videos. User evaluation suggests that users prefer our interface for learning and that our interface is effective in helping them browse or search through lecture videos.*

- ▷ Paper: Monserrat, T. J. K. P., Zhao, S., McGee, K., & Pandey, A. V. (2013, April). NoteVideo: facilitating navigation of blackboard-style lecture videos. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 1139-1148). ACM.

### [Download paper](#)

*Abstract: Khan Academy's pre-recorded blackboard-style lecture videos attract millions of online users every month. However, current video navigation tools do not adequately support the kinds of goals that students typically have, like quickly finding a particular concept in a blackboard-style lecture video. This paper reports on the development and evaluation of the new NoteVideo and its improved version, NoteVideo+, systems for identifying the conceptual 'objects' of a blackboard-based video – and then creating a summarized image of the video and using it as an in-scene navigation interface that allows users to directly jump to the video frame where that object first appeared instead of navigating it linearly through*



time. The research consisted of iteratively implementing the system and then having users perform four different navigation tasks using three different interfaces: Scrubbing, Transcript, and NoteVideo. Results of the study show that participants perform significantly better on all four tasks while using the NoteVideo and its improved version - NoteVideo+ - as compared to others

#### 4.2. Enhancing Educational Videos

☒ **Presented by: André Graën**

- ▷ Paper: Pavel, A., Reed, C., Hartmann, B., & Agrawala, M. (2014, October). Video digests: a browsable, skimmable format for informational lecture videos. In UIST (pp. 573-582).

[Download paper](#)

*Abstract: Increasingly, authors are publishing long informational talks, lectures, and distance-learning videos online. However, it is difficult to browse and skim the content of such videos using current timeline-based video players. Video digests are a new format for informational videos that afford browsing and skimming by segmenting videos into a chapter/section structure and providing short text summaries and thumbnails for each section. Viewers can navigate by reading the summaries and clicking on sections to access the corresponding point in the video. We present a set of tools to help authors create such digests using transcript-based interactions. With our tools, authors can manually create a video digest from scratch, or they can automatically generate a digest by applying a combination of algorithmic and crowdsourcing techniques and then manually refine it as needed. Feedback from first-time users suggests that our transcript-based authoring tools and automated techniques greatly facilitate video digest creation. In an evaluative crowdsourced study we find that given a short viewing time, video digests support browsing and skimming better than timeline-based or transcript-based video players.*

- ▷ Paper: Kim, J., Guo, P. J., Cai, C. J., Li, S. W. D., Gajos, K. Z., & Miller, R. C. (2014, October). Data-driven interaction techniques for improving navigation of educational videos. In Proceedings of the 27th annual ACM symposium on User interface software and technology (pp. 563-572). ACM.

[Download paper](#)

*Abstract: With an unprecedented scale of learners watching educational videos on online platforms such as MOOCs and YouTube, there is an opportunity to incorporate data generated from their interactions into the design of novel video interaction techniques. Interaction data has the potential to help not only instructors to improve their videos, but also to enrich the learning experience of educational video watchers. This paper explores the design space of data-driven interaction techniques for educational video navigation. We introduce a set of techniques that augment existing video interface widgets, including: a 2D video timeline with an embedded visualization of collective navigation traces; dynamic and non-linear timeline scrubbing; data-enhanced transcript search and keyword summary; automatic display of relevant still frames next to the video; and a visual summary representing points with high learner activity. To evaluate the feasibility of the techniques, we ran a laboratory user study with simulated learning tasks. Participants rated watching lecture videos with interaction data to be efficient and useful in completing the tasks. However, no significant differences were found in task performance, suggesting that interaction data may not always align with moment-by-moment information needs during the tasks*

#### 4.3. Survey: Human Activity Recognition from Video

□ **Topic not assigned**

- ▷ Paper: L Onofri, P Soda, M Pechenizkiy, G Iannello: A survey on using domain and contextual knowledge for human activity recognition in video streams. Expert Systems with Applications, Volume 63, 30 November 2016, Pages 97-111. Elsevier

[Download paper](#)

*Abstract: Human activity recognition has gained an increasing relevance in computer vision and it can be tackled with either non-hierarchical or hierarchical approaches. The former, also known as single-layered approaches, are those that represent and recognize human activities directly from the extracted descriptors, building a model that distinguishes among the activities contained in the training data. The latter represent and recognize human activities in terms of subevents, which are usually recognized by means of single-layered approaches. Alongside of non-hierarchical and hierarchical approaches, we observe that methods incorporating a priori knowledge and context information on the activity are getting growing interest within the community. In this work we refer to this emerging trend in computer vision as knowledge-based human activity recognition with the objective to cover the lack of a summary of these methodologies. More specifically, we survey methods and techniques used in the literature to represent and integrate knowledge and reasoning into the recognition process. We categorize them as statistical approaches, syntactic approaches and description-based approaches. In addition, we further discuss public and private datasets used in this field to promote their use and to enable the interest readers in finding useful resources. This review ends proposing main future research directions in this field.*

#### 4.4. Event Recognition in Unconstrained Videos

□ **Topic not assigned**

- ▷ Paper: Jiang, YG., Bhattacharya, S., Chang, SF. et al.: High-level event recognition in unconstrained videos. *Int J Multimed Info Retr* (2013) 2: 73.

[Download paper](#)

*Abstract: The goal of high-level event recognition is to automatically detect complex high-level events in a given video sequence. This is a difficult task especially when videos are captured under unconstrained conditions by non-professionals. Such videos depicting complex events have limited quality control, and therefore, may include severe camera motion, poor lighting, heavy background clutter, and occlusion. However, due to the fast growing popularity of such videos, especially on the Web, solutions to this problem are in high demands and have attracted great interest from researchers. In this paper, we review current technologies for complex event recognition in unconstrained videos. While the existing solutions vary, we identify common key modules and provide detailed descriptions along with some insights for each of them, including extraction and representation of low-level features across different modalities, classification strategies, fusion techniques, etc. Publicly available benchmark datasets, performance metrics, and related research forums are also described. Finally, we discuss promising directions for future research.*

#### 4.5. Multimodal Information for Event Summarization and Sentiment Analysis

□ **Topic not assigned**

- ▷ Paper: Rajiv Ratn Shah, Yi Yu, Akshay Verma, Suhua Tang, Anwar Dilawar Shaikh, Roger Zimmermann, Leveraging multimodal information for event summarization and concept-level sentiment analysis, In *Knowledge-Based Systems*, Volume 108, 2016, Pages 102-109, ISSN 0950-7051

[Download paper](#)

*Abstract: The rapid growth in the amount of user-generated content (UGC)s online necessitates for social media companies to automatically extract knowledge structures (concepts) from photos and videos to provide diverse multimedia-related services. However, real-world photos and videos are complex and noisy, and extracting semantics and santics from the multimedia content alone is a very difficult task because suitable concepts may be exhibited in different representations. Hence, it is desirable to analyze UGCs from multiple modalities for a better understanding. To this end, we first present the EventBuilder system that deals with semantics understanding and automatically generates a multimedia summary for a given event in real-time by leveraging different social media such as Wikipedia and Flickr. Subsequently, we present the EventSensor system that aims to address santics understanding and produces a multimedia summary for a given mood. It extracts concepts and mood tags from visual content and textual metadata of UGCs, and exploits them in supporting several significant multimedia-related services such as a musical multimedia summary. Moreover, EventSensor supports santics-based event summarization by leveraging EventBuilder as its semantics engine component. Experimental results confirm that both EventBuilder and EventSensor outperform their baselines and efficiently summarize knowledge structures on the YFCC100M dataset.*

#### 4.6. Video Representation for Action Recognition

□ **Topic not assigned**

- ▷ Paper: Wang, H., Oneata, D., Verbeek, J. et al.: A Robust and Efficient Video Representation for Action Recognition *Int J Comput Vis* (2016) 119: 219.

[Download paper](#)

*Abstract: This paper introduces a state-of-the-art video representation and applies it to efficient action recognition and detection. We first propose to improve the popular dense trajectory features by explicit camera motion estimation. More specifically, we extract feature point matches between frames using SURF descriptors and dense optical flow. The matches are used to estimate a homography with RANSAC. To improve the robustness of homography estimation, a human detector is employed to remove outlier matches from the human body as human motion is not constrained by the camera. Trajectories consistent with the homography are considered as due to camera motion, and thus removed. We also use the homography to cancel out camera motion from the optical flow. This results in significant improvement on motion-based HOF and MBH descriptors. We further explore the recent Fisher vector as an alternative feature encoding approach to the standard bag-of-words (BOW) histogram, and consider different ways to include spatial layout information in these encodings. We present a large and varied set of evaluations, considering (i) classification of short basic actions on six datasets, (ii) localization of such actions in feature-length movies, and (iii) large-scale recognition of complex events. We find that our improved trajectory features significantly outperform previous dense trajectories, and that Fisher vectors are superior to BOW encodings for video recognition tasks. In all three tasks, we show substantial improvements over the state-of-the-art results.*

## 5. Social and Explanatory Behaviour

### 5.1. Grand Challenges of Science Robotics

☒ **Presented by: David Löwens**

- ▷ Paper: Yang, Guang-Zhong, Jim Bellingham, Pierre E. Dupont, Peer Fischer, Luciano Floridi, Robert Full, Neil Jacobstein et al. "The grand challenges of Science Robotics." *Science robotics* 3, no. 14 (2018): eaar7650.

[Download paper](#)

*Abstract: One of the ambitions of Science Robotics is to deeply root robotics research in science while developing novel robotic platforms that will enable new scientific discoveries. Of our 10 grand challenges, the first 7 represent underpinning technologies that have a wider impact on all application areas of robotics. For the next two challenges, we have included social robotics and medical robotics as application-specific areas of development to highlight the substantial societal and health impacts that they will bring. Finally, the last challenge is related to responsible innovation and how ethics and security should be carefully considered as we develop the technology further.*

### 5.2. Context-Aware Analysis and Annotation of Human-Agent Interaction

☒ **Presented by: Isabel Laqua**

- ▷ Paper: Baur, Tobias, Gregor Mehlmann, Ionut Damian, Florian Lingenfeller, Johannes Wagner, Birgit Lugin, Elisabeth André, and Patrick Gebhard. "Context-Aware Automated Analysis and Annotation of Social Human-Agent Interactions." *ACM Transactions on Interactive Intelligent Systems (TiIS)* 5, no. 2 (2015): 11.

[Download paper](#)

*Abstract: The outcome of interpersonal interactions depends not only on the contents that we communicate verbally, but also on nonverbal social signals. As a lack of social skills is a common problem for a significant number of people, serious games and other training environments have recently become the focus of research. In this work we present NovA (Nonverbal behavior Analyzer), a system that analyzes and facilitates the interpretation of social signals automatically in a bi-directional interaction with a conversational agent. It records data of interactions, detects relevant social cues, and creates descriptive statistics for the recorded data with respect to the agents behavior and the context of the situation. This enhances the possibilities for researchers to automatically label corpora of human-agent interactions and to give users feedback on strengths and weaknesses of their social behavior*

### 5.3. Machine Learning for Semi-Automated Annotation of Social Signals

☐ **Topic not assigned**

- ▷ Paper: Baur, Tobias, Ionut Damian, Florian Lingenfeller, Johannes Wagner, and Elisabeth André. "Nova: Automated analysis of nonverbal signals in social interactions." In *International Workshop on Human Behavior Understanding*, pp. 160-171. Springer, Cham, 2013.

[Download paper](#)

*Abstract: Previous studies have shown that the success of interpersonal interaction depends not only on the contents we communicate explicitly, but also on the social signals that are conveyed implicitly. In this paper, we present NovA (NONVerbal behavior Analyzer), a system that analyzes and facilitates the interpretation of social signals conveyed by gestures, facial expressions and others automatically as a basis for computer-enhanced social coaching. NovA records data of human interactions, automatically detects relevant behavioral cues as a measurement for the quality of an interaction and creates descriptive statistics for the recorded data. This enables us to give a user online generated feedback on strengths and weaknesses concerning his social behavior, as well as elaborate tools for online analysis and annotation.*

- ▷ Paper: Wagner, Johannes, Tobias Baur, Dominik Schiller, Yue Zhang, Björn Schuller, Michel Valstar, and Elisabeth André. "Show Me What You've Learned: Applying Cooperative Machine Learning for the Semi-Automated Annotation of Social Signals." *XAI 2018*: 171.

[Download paper](#)

*Abstract: In this paper we suggest the use of Cooperative Machine Learning (CML) to reduce manual labelling efforts while simultaneously generating an intuitive understanding of the learning process of a classification system. To this end, we introduce the open-source tool NOVA, which aims to combine human intelligence and machine learning to annotate social signals in large multi-modal corpora. NOVA features a semi-automated labelling process in which users are provided with immediate visual feedback on the predictions, which affords insights into the strengths and weaknesses of the underlying classification system. Following an interactive and exploratory workflow, the performance of the model can be improved by manual revision of the predictions, a process that uses confidence values to guide the inspection.*

### 5.4. Robust Interactive Learning Social Robot

☒ **Presented by: Aysegül Karabag**

- ▷ Paper: de Jong, Michiel, Kevin Zhang, Aaron M. Roth, Travers Rhodes, Robin Schmucker, Chenghui Zhou, Sofia Ferreira, João Cartucho, and Manuela Veloso. "Towards a robust in-

teractive and learning social robot." In Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, pp. 883-891. International Foundation for Autonomous Agents and Multiagent Systems, 2018.

[Download paper](#)

*Abstract: Pepper is a humanoid robot, specifically designed for social interaction, that has been deployed in a variety of public environments. A programmable version of Pepper is also available, enabling our focused research on perception and behavior robustness and capabilities of an interactive social robot. We address Pepper perception by integrating state-of-the-art vision and speech recognition systems and experimentally analyzing their effectiveness. As we recognize limitations of the individual perceptual modalities, we introduce a multi-modality approach to increase the robustness of human social interaction with the robot. We combine vision, gesture, speech, and input from an onboard tablet, a remote mobile phone, and external microphones. Our approach includes the proactive seeking of input from a different modality, adding robustness to the failures of the separate components. We also introduce a learning algorithm to improve communication capabilities over time, updating speech recognition through social interactions. Finally, we realize the rich robot body-sensory data and introduce both a nearest-neighbor and a deep learning approach to enable Pepper to classify and speak up a variety of its own body motions. We view the contributions of our work to be relevant both to Pepper specifically and to other general social robots.*

### 5.5. Attitudes towards potentially uncanny robots

☒ **Presented by: Felix Piepenbrink**

▷ Paper: Rosenthal-von der Pütten, Astrid M., and Nicole C. Krämer. "Individuals' evaluations of and attitudes towards potentially uncanny robots." *International Journal of Social Robotics* 7, no. 5 (2015): 799-824.

[Download paper](#)

*Abstract: In the present work we provide an overview and categorization of explanatory approaches for the uncanny valley effect and present an empirical study. Against the background of the uncanny valley hypothesis, the study utilized qualitative interviews in which participants were presented with pictures and videos of potentially uncanny humanoid and android robots to explore participants' evaluations of very human-like robots, their attitudes about these robots, and their emotional reactions towards these robots. In this regard, the influence of the robots' appearance, movement and the context of HRI were examined. Results showed that, contrasting the hypothesis, participants reported not only negative, but also positive emotional reactions towards the possibly uncanny robots. The robots' appearance was of great importance for the participants, because certain characteristics were equalized with certain abilities, merely human appearance without a connected functionality was not appreciated, and human rules of attractiveness were applied to the android robots. The analysis also demonstrated the importance of the robots' movements and the social context they were placed in. The importance of two possible causes and explanations of the uncanny valley, namely uncertainty at category boundaries (cf. Ramey in Proceedings of views of the uncanny valley workshop: IEEE-RAS international conference on humanoid robots, 2005; Proceedings of the ICCS/CogSci-2006 long symposium "Toward Social Mechanisms of Android Science", 2006) and subconscious fears of being replaced (cf. MacDorman & Ishiguro in Interact Stud 7(3):297-337, 2006) were explored in this work. On this reflective level of evaluation we found some support for the assumptions that participants experienced uncertainty how to categorize android robots (as human or machine) and that some (but not all) participants felt uncomfortable at the thought to be replaced by robots.*

### 5.6. Grand Challenges for Ambient Intelligence

□ **Topic not assigned**

▷ Paper: Streitz, Norbert, Dimitris Charitos, Maurits Kaptein, and Marc Böhlen. "Grand challenges for ambient intelligence and implications for design contexts and smart societies." *Journal of Ambient Intelligence and Smart Environments* 11, no. 1 (2019): 87-107.

[Download paper](#)

*Abstract: This paper highlights selected grand challenges that concern especially the social and the design dimensions of research and development in Ambient Intelligence (AmI) and Smart Environments (SmE). Due to the increasing deployment and usage of 'smart' technologies determining a wide range of everyday life activities, there is an urgent need to reconsider their societal implications and how to address these implications with appropriate design methods. The paper presents four perspectives on the subject grounded in different approaches. First, introducing and reflecting on the implications of the 'smart-everything' paradigm, the resulting design trade-offs and their application to smart cities. Second, discussing the potential of non-verbal communication for informing the design of spatial interfaces for AmI design practices. Third, reflecting on the role of new data categories such as 'future data' and the role of uncertainty and their implications for the next generation of AmI environments. Finally, debating*

*the merits and shortfalls of the world's largest professional engineering community effort to craft a global standards body on ethically aligned design for autonomous and intelligent systems. The paper benefits from taking different perspectives on common issues, provides commonalities and relationships between them and provides anchor points for important challenges in the field of ambient intelligence.*

### 5.7. Survey Explainable Agents and Robots

□ **Topic not assigned**

- ▷ Paper: Anjomshoae, Sule, Amro Najjar, Davide Calvaresi, and Kary Främling. "Explainable agents and robots: Results from a systematic literature review." In Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, pp. 1078-1088. International Foundation for Autonomous Agents and Multiagent Systems, 2019.

[Download paper](#)

*Abstract: Humans are increasingly relying on complex systems that heavily adopts Artificial Intelligence (AI) techniques. Such systems are employed in a growing number of domains, and making them explainable is an impelling priority. Recently, the domain of explainable Artificial Intelligence (XAI) emerged with the aims of fostering transparency and trustworthiness. Several reviews have been conducted. Nevertheless, most of them deal with data-driven XAI to overcome the opaqueness of black-box algorithms. Contributions addressing goal-driven XAI (e.g., explainable agency for robots and agents) are still missing. This paper aims at filling this gap, proposing a Systematic Literature Review. The main findings are (i) a considerable portion of the papers propose conceptual studies, or lack evaluations or tackle relatively simple scenarios; (ii) almost all of the studied papers deal with robots/agents explaining their behaviors to the human users, and very few works addressed inter-robot (inter-agent) explainability. Finally, (iii) while providing explanations to non-expert users has been outlined as a necessity, only a few works addressed the issues of personalization and context-awareness.*

### 5.8. Explanation in AI: Insights from the social sciences

□ **Topic not assigned**

- ▷ Paper: Miller, T. (2017). Explanation in Artificial Intelligence: Insights from the social sciences. [arXiv:1706.07269]

[Download paper](#)

*Abstract: There has been a recent resurgence in the area of explainable artificial intelligence as researchers and practitioners seek to make their algorithms more understandable. Much of this research is focused on explicitly explaining decisions or actions to a human observer, and it should not be controversial to say that looking at how humans explain to each other can serve as a useful starting point for explanation in artificial intelligence. However, it is fair to say that most work in explainable artificial intelligence uses only the researchers' intuition of what constitutes a 'good' explanation. There exists vast and valuable bodies of research in philosophy, psychology, and cognitive science of how people define, generate, select, evaluate, and present explanations, which argues that people employ certain cognitive biases and social expectations towards the explanation process. This paper argues that the field of explainable artificial intelligence should build on this existing research, and reviews relevant papers from philosophy, cognitive psychology/science, and social psychology, which study these topics. It draws out some important findings, and discusses ways that these can be infused with work on explainable artificial intelligence.*

### 5.9. Contextual Utility Affects Perceived Quality of Explanations

□ **Topic not assigned**

- ▷ Paper: Vasilyeva, Nadya, Daniel Wilkenfeld, and Tania Lombrozo. "Contextual utility affects the perceived quality of explanations." *Psychonomic bulletin & review* 24, no. 5 (2017): 1436-1450.

[Download paper](#)

*Abstract: Are explanations of different kinds (formal, mechanistic, teleological) judged differently depending on their contextual utility, defined as the extent to which they support the kinds of inferences required for a given task? We report three studies demonstrating that the perceived "goodness" of an explanation depends on the evaluator's current task: Explanations receive a relative boost when they support task-relevant inferences, even when all three explanation types are warranted. For example, mechanistic explanations receive higher ratings when participants anticipate making further inferences on the basis of proximate causes than when they anticipate making further inferences on the basis of category membership or functions. These findings shed light on the functions of explanation and support pragmatic and pluralist approaches to explanation.*

### 5.10. Trends and Trajectories for Explainable Systems

□ **Topic not assigned**

- ▷ Paper: Abdul, Ashraf, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. "Trends and trajectories for explainable, accountable and intelligible systems: An HCI research

agenda." In Proceedings of the 2018 CHI conference on human factors in computing systems, p. 582. ACM, 2018.

[Download paper](#)

Abstract: *Advances in artificial intelligence, sensors and big data management have far-reaching societal impacts. As these systems augment our everyday lives, it becomes increasingly important for people to understand them and remain in control. We investigate how HCI researchers can help to develop accountable systems by performing a literature analysis of 289 core papers on explanations and explainable systems, as well as 12,412 citing papers. Using topic modeling, co-occurrence and network analysis, we mapped the research space from diverse domains, such as algorithmic accountability, interpretable machine learning, context-awareness, cognitive psychology, and software learnability. We reveal fading and burgeoning trends in explainable systems, and identify domains that are closely connected or mostly isolated. The time is ripe for the HCI community to ensure that the powerful new autonomous systems have intelligible interfaces built-in. From our results, we propose several implications and directions for future research towards this goal.*

#### 5.11. What does Explainable AI Really Mean?

**Topic not assigned**

▷ Paper: Doran, Derek, Sarah Schulz, and Tarek R. Besold. "What does explainable AI really mean? A new conceptualization of perspectives." arXiv preprint arXiv:1710.00794 (2017).

[Download paper](#)

Abstract: *We characterize three notions of explainable AI that cut across research fields: opaque systems that offer no insight into its algorithmic mechanisms; interpretable systems where users can mathematically analyze its algorithmic mechanisms; and comprehensible systems that emit symbols enabling user-driven explanations of how a conclusion is reached. The paper is motivated by a corpus analysis of NIPS, ACL, COGSCI, and ICCV/ECCV paper titles showing differences in how work on explainable AI is positioned in various fields. We close by introducing a fourth notion: truly explainable systems, where automated reasoning is central to output crafted explanations without requiring human post processing as final step of the generative process.*

#### 5.12. Explain or Not – Music Recommendations

**Presented by: Vedat Latifi**

▷ Paper: Millemcamp, Martijn, Nyi Nyi Htun, Cristina Conati, and Katrien Verbert. "To explain or not to explain: the effects of personal characteristics when explaining music recommendations." In IUI, pp. 397-407. 2019.

[Download paper](#)

Abstract: *Recommender systems have been increasingly used in online services that we consume daily, such as Facebook, Netflix, YouTube, and Spotify. However, these systems are often presented to users as a "black box", i.e. the rationale for providing individual recommendations remains unexplained to users. In recent years, various attempts have been made to address this black box issue by providing textual explanations or interactive visualisations that enable users to explore the provenance of recommendations, and benefits in terms of precision and user satisfaction, among others, have been demonstrated. Previous research had also indicated that personal characteristics such as domain knowledge, trust propensity and persistence may also play an important role on such perceived benefits. Yet, to date, little is known about the effects of personal characteristics when explaining recommendations. To address this gap, we developed a music recommender system with explanations and conducted an online study using a within-subject design. We captured various personal characteristics of participants and administered both qualitative and quantitative evaluation methods. Results indicated that personal characteristics have some significant influence on the interaction and perception of recommender systems and that this influence changes by adding explanations. Especially people with a low need for cognition benefited from explained recommendations. For people with a high need for cognition, we observed that explanations could lower their confidence. Based on these results, we present some first design implications for explaining recommendations.*

#### 5.13. Rigorous Science of Interpretable Machine Learning

**Topic not assigned**

▷ Paper: Doshi-Velez, Finale, and Been Kim. "Towards a rigorous science of interpretable machine learning." arXiv preprint arXiv:1702.08608 (2017).

[Download paper](#)

Abstract: *As machine learning systems become ubiquitous, there has been a surge of interest in interpretable machine learning: systems that provide explanation for their outputs. These explanations are often used to qualitatively assess other criteria such as safety or non-discrimination. However, despite the interest in interpretability, there is very little consensus on what interpretable machine learning is and how it should be measured. In this position paper, we first define interpretability and describe when*

interpretability is needed (and when it is not). Next, we suggest a taxonomy for rigorous evaluation and expose open questions towards a more rigorous science of interpretable machine learning.

#### 5.14. Dialogical Models of Explanation

□ **Topic not assigned**

▷ Paper: Walton, Douglas. "Dialogical Models of Explanation." *ExaCt* 2007 (2007): 1-9.

[Download paper](#)

Abstract: *This paper takes on the task of providing a formal system of dialogue CE in which the speech acts of requesting and providing an explanation are represented as dialogue moves in the system. CE has opening rules, locution rules, dialogue rules, success rules and closing rules. The system is meant to be simple and basic, to provide a platform for developing more specialized formal dialogue systems of explanation used for specific purposes. The dialogical theory of explanation postulates that an explanation is a dialogue between two parties, one of whom asks a question requesting understanding of something which he or she claims not to understand, while the other offers a response that claims to convey the requested understanding to the party asking the question.*

### 6. Ethical Computing

#### 6.1. Review Future and Ethical Perspectives on AI

□ **Topic not assigned**

▷ Paper: Torresen, Jim. "A review of future and ethical perspectives of robotics and AI." *Frontiers in Robotics and AI* 4 (2018): 75.

[Download paper](#)

Abstract: *In recent years, there has been increased attention on the possible impact of future robotics and AI systems. Prominent thinkers have publicly warned about the risk of a dystopian future when the complexity of these systems progresses further. These warnings stand in contrast to the current state-of-the-art of the robotics and AI technology. This article reviews work considering both the future potential of robotics and AI systems, and ethical considerations that need to be taken in order to avoid a dystopian future. References to recent initiatives to outline ethical guidelines for both the design of systems and how they should operate are included.*

#### 6.2. Designing AI for Human Values

☒ **Presented by: Vijona Krasniqi**

▷ Paper: Dignum, Virginia. "Responsible Artificial Intelligence: Designing AI for Human Values." (2017).

[Download paper](#)

Abstract: *Artificial intelligence (AI) is increasingly affecting our lives in smaller or greater ways. In order to ensure that systems will uphold human values, design methods are needed that incorporate ethical principles and address societal concerns. In this paper, we explore the impact of AI in the case of the expected effects on the European labor market, and propose the accountability, responsibility and transparency (ART) design principles for the development of AI systems that are sensitive to human values.*

#### 6.3. Ethical Considerations in AI Courses

□ **Topic not assigned**

▷ Paper: Burton, Emanuelle, Judy Goldsmith, Sven Koenig, Benjamin Kuipers, Nicholas Mattei, and Toby Walsh. "Ethical considerations in artificial intelligence courses." *AI magazine* 38, no. 2 (2017): 22-34.

[Download paper](#)

Abstract: *The recent surge in interest in ethics in artificial intelligence (AI) may leave many educators wondering how to address moral, ethical, and philosophical issues in their AI courses. As instructors we want to develop curriculum that not only prepares students to be AI practitioners, but also to understand the moral, ethical, and philosophical impacts that AI will have on society. In this article we provide practical case studies and links to resources for use by AI educators. We also provide concrete suggestions on how to integrate AI ethics into a general AI course and how to teach a stand-alone AI ethics course.*

#### 6.4. Teaching Computer Ethics through Science Fiction

☒ **Presented by: Kreshnik Duci**

▷ Paper: Burton, Emanuelle, Judy Goldsmith, and Nicholas Mattei. "How to teach computer ethics through science fiction." *Communications of the ACM* 61, no. 8 (2018): 54-64.

[Download paper](#)

Abstract: *Computer science faculty have a responsibility to teach students to recognize both the larger ethical issues and particular responsibilities that are part and parcel of their work as technologists. This is, however, a kind of teaching for which most of us have not been trained, and that faculty and students approach with some trepidation. In this article, we explore the use of science fiction as a tool to enable those*

teaching artificial intelligence to engage students and practitioners about the scope and implications of current and future work in computer science. We have spent several years developing a creative approach to teaching computer ethics, through a course we call “Science Fiction and Computer Ethics.”

#### 6.5. Research Priorities for Robust and Beneficial AI

☒ **Presented by: Merouan Abou-Abdo**

▷ Paper: Russell, Stuart, Daniel Dewey, and Max Tegmark. “Research priorities for robust and beneficial artificial intelligence.” *Ai Magazine* 36, no. 4 (2015): 105-114.

[Download paper](#)

Abstract: *Success in the quest for artificial intelligence has the potential to bring unprecedented benefits to humanity, and it is therefore worthwhile to investigate how to maximize these benefits while avoiding potential pitfalls. This article gives numerous examples (which should by no means be construed as an exhaustive list) of such worthwhile research aimed at ensuring that AI remains robust and beneficial.*

#### 6.6. Classification of Moral Mediation by Ubiquitous Machines

☐ **Topic not assigned**

▷ Paper: Van de Voort, Marlies, Wolter Pieters, and Luca Consoli. “Refining the ethics of computer-made decisions: a classification of moral mediation by ubiquitous machines.” *Ethics and Information Technology* 17, no. 1 (2015): 41-56.

[Download paper](#)

Abstract: *In the past decades, computers have become more and more involved in society by the rise of ubiquitous systems, increasing the number of interactions between humans and IT systems. At the same time, the technology itself is getting more complex, enabling devices to act in a way that previously only humans could, based on developments in the fields of both robotics and artificial intelligence. This results in a situation in which many autonomous, intelligent and context-aware systems are involved in decisions that affect their environment. These relations between people, machines, and decisions can take many different forms, but thus far, a systematic account of machine-assisted moral decisions is lacking. This paper investigates the concept of machine-assisted moral decisions from the perspective of technological mediation. It is argued that modern machines do not only have morality in the sense of mediating the actions of humans, but that, by making their own decisions within their relations with humans, mediate morality itself. A classification is proposed to differentiate between four different types of moral relations. The moral aspects within the decisions these systems make are combined into three dimensions that describe the distinct characteristics of different types of moral mediation by machines. Based on this classification, specific guidelines for moral behavior can be provided for these systems.*

#### 6.7. Ethics of Healthcare Robotics

☒ **Presented by: Anna Ushakov**

▷ Paper: Stahl, Bernd Carsten, and Mark Coeckelbergh. “Ethics of healthcare robotics: Towards responsible research and innovation.” *Robotics and Autonomous Systems* 86 (2016): 152-161.

[Download paper](#)

Abstract: *How can we best identify, understand, and deal with ethical and societal issues raised by healthcare robotics? This paper argues that next to ethical analysis, classic technology assessment, and philosophical speculation we need forms of reflection, dialogue, and experiment that come, quite literally, much closer to innovation practices and contexts of use. The authors discuss a number of ways how to achieve that. Informed by their experience with “embedded” ethics in technical projects and with various tools and methods of responsible research and innovation, the paper identifies “internal” and “external” forms of dialogical research and innovation, reflections on the possibilities and limitations of these forms of ethical–technological innovation, and explores a number of ways how they can be supported by policy at national and supranational level.*

#### 6.8. Social Dilemma of Autonomous Vehicles

☒ **Presented by: Niklas Pick**

▷ Paper: Bonnefon, Jean-François, Azim Shariff, and Iyad Rahwan. “The social dilemma of autonomous vehicles.” *Science* 352, no. 6293 (2016): 1573-1576.

[Download paper](#)

Abstract: *Autonomous vehicles (AVs) should reduce traffic accidents, but they will sometimes have to choose between two evils, such as running over pedestrians or sacrificing themselves and their passenger to save the pedestrians. Defining the algorithms that will help AVs make these moral decisions is a formidable challenge. We found that participants in six Amazon Mechanical Turk studies approved of utilitarian AVs (that is, AVs that sacrifice their passengers for the greater good) and would like others to buy them, but they would themselves prefer to ride in AVs that protect their passengers at all costs. The study participants disapprove of enforcing utilitarian regulations for AVs and would be less willing to buy such an AV. Accordingly, regulating for utilitarian algorithms may paradoxically increase casualties by postponing the adoption of a safer technology.*



## 6.9. Fair, Transparent, Accountable Decision Making

□ **Topic not assigned**

- ▷ Paper: Lepri, Bruno, Nuria Oliver, Emmanuel Letouzé, Alex Pentland, and Patrick Vinck. "Fair, transparent, and accountable algorithmic decision-making processes." *Philosophy & Technology* 31, no. 4 (2018): 611-627.

[Download paper](#)

*Abstract: The combination of increased availability of large amounts of fine-grained human behavioral data and advances in machine learning is presiding over a growing reliance on algorithms to address complex societal problems. Algorithmic decision-making processes might lead to more objective and thus potentially fairer decisions than those made by humans who may be influenced by greed, prejudice, fatigue, or hunger. However, algorithmic decision-making has been criticized for its potential to enhance discrimination, information and power asymmetry, and opacity. In this paper, we provide an overview of available technical solutions to enhance fairness, accountability, and transparency in algorithmic decision-making. We also highlight the criticality and urgency to engage multi-disciplinary teams of researchers, practitioners, policy-makers, and citizens to co-develop, deploy, and evaluate in the real-world algorithmic decision-making processes designed to maximize fairness and transparency. In doing so, we describe the Open Algorithms (OPAL) project as a step towards realizing the vision of a world where data and algorithms are used as lenses and levers in support of democracy and development.*

## 6.10. AI Fairness 360 – Toolkit Algorithmic Bias

□ **Topic not assigned**

- ▷ Paper: Bellamy, Rachel KE, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia et al. "AI fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias." arXiv preprint arXiv:1810.01943 (2018).

[Download paper](#)

*Abstract: Fairness is an increasingly important concern as machine learning models are used to support decision making in high-stakes applications such as mortgage lending, hiring, and prison sentencing. This paper introduces a new open source Python toolkit for algorithmic fairness, AI Fairness 360 (AIF360), released under an Apache v2.0 license. The main objectives of this toolkit are to help facilitate the transition of fairness research algorithms to use in an industrial setting and to provide a common framework for fairness researchers to share and evaluate algorithms. The package includes a comprehensive set of fairness metrics for datasets and models, explanations for these metrics, and algorithms to mitigate bias in datasets and models. It also includes an interactive Web experience (this https URL) that provides a gentle introduction to the concepts and capabilities for line-of-business users, as well as extensive documentation, usage guidance, and industry-specific tutorials to enable data scientists and practitioners to incorporate the most appropriate tool for their problem into their work products. The architecture of the package has been engineered to conform to a standard paradigm used in data science, thereby further improving usability for practitioners. Such architectural design and abstractions enable researchers and developers to extend the toolkit with their new algorithms and improvements, and to use it for performance benchmarking. A built-in testing infrastructure maintains code quality.*

## 7. Historic Milestones

### 7.1. Sword of Damocles

□ **Topic not assigned**

- ▷ Paper: Sutherland, Ivan E. A head-mounted three dimensional display. Proceedings of the 1968 Fall Joint Computer Conference, San Francisco, CA, December 1968, Pages 757-764. ACM

[Download paper](#)

*Abstract: The fundamental idea behind the three-dimensional display is to present the user with a perspective image which changes as he moves. The retinal image of the real objects which we see is, after all, only two-dimensional. Thus if we can place suitable two-dimensional images on the observer's retinas, we can create the illusion that he is seeing a three-dimensional object. Although stereo presentation is important to the three-dimensional illusion, it is less important than the change that takes place in the image when the observer moves his head. The image presented by the three-dimensional display must change in exactly the way that the image of a real object would change for similar motions of the user's head. Psychologists have long known that moving perspective images appear strikingly three-dimensional even without stereo presentation; the three-dimensional display described in this paper depends heavily on this "kinetic depth effect."*

## 7.2. Augmenting Human Intellect

**Topic not assigned**

- ▷ Paper: C Engelbart, and WK English: A research center for augmenting human intellect. Proceedings of the 1968 Fall Joint Computer Conference, San Francisco, CA, December 1968, Vol. 33, pp. 395-410. ACM

[Download paper](#)

Abstract: *This paper describes a multisponsor research center at Stanford Research Institute in man-computer interaction.*

## 7.3. Put That There

**Topic not assigned**

- ▷ Paper: RA Bolt: "Put-that-there": Voice and gesture at the graphics interface. SIGGRAPH '80 Proceedings of the 7th annual conference on Computer graphics and interactive techniques, Pages 262-270, 1980.

[Download paper](#)

Abstract: *Recent technological advances in connected-speech recognition and position sensing in space have encouraged the notion that voice and gesture inputs at the graphics interface can converge to provide a concerted, natural user modality. The work described herein involves the user commanding simple shapes about a large-screen graphics display surface. Because voice can be augmented with simultaneous pointing, the free usage of pronouns becomes possible, with a corresponding gain in naturalness and economy of expression. Conversely, gesture aided by voice gains precision in its power to reference.*

## 7.4. Tangible Bits

**Presented by: René Springer**

- ▷ Paper: Ishii, Hiroshi, and Brygg Ullmer. "Tangible bits: towards seamless interfaces between people, bits and atoms." Proceedings of the ACM SIGCHI Conference on Human factors in computing systems. ACM, 1997.

[Download paper](#)

Abstract: *This paper presents our vision of Human Computer Interaction (HCI): "Tangible Bits." Tangible Bits allows users to "grasp & manipulate" bits in the center of users' attention by coupling the bits with everyday physical objects and architectural surfaces. Tangible Bits also enables users to be aware of background bits at the periphery of human perception using ambient display media such as light, sound, airflow, and water movement in an augmented space. The goal of Tangible Bits is to bridge the gaps between both cyberspace and the physical environment, as well as the foreground and background of human activities. This paper describes three key concepts of Tangible Bits: interactive surfaces; the coupling of bits with graspable physical objects; and ambient media for background awareness. We illustrate these concepts with three prototype systems – the metaDESK, transBOARD and ambientROOM – to identify underlying research issues.*

## 8. Transdisciplinary Issues

### 8.1. Systemic-Functional Theory of Language: Semantic Networks

**Topic not assigned**

- ▷ Paper: R Hasan, C Cloran, G Williams and A Lukin: Semantic networks: the description of linguistic meaning in SFL. In R. Hasan, C. MIM. Matthiessen & J. J. Webster (eds), Continuing Discourse on Language: A Functional Perspective Volume 2, 1 edn, pp. 697-738, Equinox Publishing Ltd, London, UK and Oakville, USA, 2005.

[Download paper](#)

Abstract: *In SFL, the semantics of a language calls for as much attention as its lexicogrammar: in fact, meaning and wording are two sides of the same coin; the description of both is equally central to understanding 'how language works' (Halliday, McIntosh and Strevens, 1964) – which has been Halliday's agenda since the beginning of his engagement with linguistics. But what actually led SFL into the exploration of semantics as a legitimate domain for description was not these theoretical considerations, per se; rather, like other aspects of the evolution of SFL, interest in semantics too arose in attempts to resolve certain problems in the course of research during the 1960s. This chapter presents one perspective on the course of this development, specifically with respect to semantic networks as a resource for the analysis of meaning.*

### 8.2. Systemic-Functional Theory of Language: Academic Writing

**Topic not assigned**

- ▷ Paper: D Motta-Roth: The Role of Context in Academic Text Production and Writing Pedagogy. In: C Bazerman, A Bonini, D Figueiredo (eds), Genre in a Changing World, pp. 317-336, Parlor Press, 2009.

[Download paper](#)

**Abstract:** *The problem of text production in academic genres has been a challenge for undergraduate and graduate students as well as for writing teachers from different departments. Previous research has provided important results on the structural aspects of academic genres (Swales, 1990) and the discursive construction of identity in academic writing (Ivanic, 1998). However, few studies have concentrated on the contributions of exploration and reflection on context to actual teaching practices. From the perspective of Systemic-Functional Linguistics (SFL), in this paper I would like to focus on the reciprocal relationship between text and context, i.e., the way context can be recreated by analysis of text and vice versa. The aim is to point out some practical implications derived from the use of SFL principles in academic writing teaching and research through context exploration. The focus will be on writing activities that aim at fostering students' awareness about the connections between contextual features (activity, identity, relations as well as the role performed by text in the situation) and their respective linguistic realizations (expression of content, instantiation of relationships between interlocutors, and organization of text). One of the main challenges in language education and research is to teach creative ways to negotiate the norms of the language system (grammar) within the academic culture: the set of meanings, rules, values, power relations and relevant genres that constitute the social practices of a community. Educating students about the uses of language in specific contexts depends on clear descriptions of the connections between text and context. Public discourse on academic publication in Brazil is mainly issued by the Ministry of Education through its two main Research Funding Agencies, CAPES and CNPq, which hold quantitative and qualitative expectations about scholars' intellectual production but offer no substantial line of financial support for pedagogic research and course development. Thus full-fledged writing programs are seldom found in Brazilian universities. Very often what we find is some individual or collective teaching initiatives that have survived defying all the odds (e.g., lack of personnel and financial resources) situated in specific institutions.*

### 8.3. Activity Theory: Understanding Acting in Complex Environments

□ **Topic not assigned**

- ▷ Paper: Leena Norros: "Understanding Acting in Complex Environments: Building a Synergy of Cultural-Historical Activity Theory, Peirce, and Ecofunctionalism". In: *Mind, Culture, and Activity*, 25:1, 68-85, 2018

[Download paper](#)

**Abstract:** *A method of activity analysis is proposed that exploits synergy among cultural-historical activity theory, Peircean sign theory, and ecofunctionalism in response to challenges of modern work. The method comprises an operationalisation of the object of activity via the concept of core task, which enables contextual description of actions. On this basis, a semiotic analysis can be accomplished that reveals generic habitual ways of acting. The variation found in their interpretive strength for comprehending situations predicts differences in the mastery of work and learning in work. The method also allows definition of systems usability as a holistic quality-linked concept for purposes of artefact design.*

### 8.4. Activity Theory: Plans as Situated Action

□ **Topic not assigned**

- ▷ Paper: Bardram, Jakob E. "Plans as situated action: an activity theory approach to workflow systems." *Proceedings of the Fifth European Conference on Computer Supported Cooperative Work*. Springer Netherlands, 1997.

[Download paper](#)

**Abstract:** *Within the community of CSCW the notion and nature of workflow systems as prescriptions of human work has been debated and criticised. Based on the work of Suchman (1987) the notion of situated action has often been viewed as opposed to planning work. Plans, however, do play an essential role in realising work. Based on experiences from designing a computer system that supports the collaboration within a hospital, this paper discusses how plans themselves are made out of situated action, and in return are realised in situ. Thus, work can be characterised as situated planning. This understanding is backed up by Activity Theory, which emphasises the connection between plans and the contextual conditions for realising these plans in actual work.*

#### Source

Texts can be downloaded from the [Learnweb-course](#).

List of the 38 topics which have not been assigned yet: 1.1., 1.3., 1.4., 2.1., 2.2., 2.3., 2.4., 2.9., 3.3., 3.4., 3.5., 3.6., 4.1., 4.3., 4.4., 4.5., 4.6., 5.3., 5.6., 5.7., 5.8., 5.9., 5.10., 5.11., 5.13., 5.14., 6.1., 6.3., 6.6., 6.9., 6.10., 7.1., 7.2., 7.3., 8.1., 8.2., 8.3., 8.4., Please see individual papers for a description.