



Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

The Importance of Context When Recommending TV Content

Dataset and Algorithms

Kristoffersen, Miklas Strøm; Shepstone, Sven Ewan; Tan, Zheng-Hua

Published in:

I E E E Transactions on Multimedia

DOI (link to publication from Publisher):

[10.1109/TMM.2019.2944214](https://doi.org/10.1109/TMM.2019.2944214)

Creative Commons License

CC BY 4.0

Publication date:

2020

Document Version

Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Kristoffersen, M. S., Shepstone, S. E., & Tan, Z-H. (2020). The Importance of Context When Recommending TV Content: Dataset and Algorithms. *I E E E Transactions on Multimedia*, 22(6), 1531-1541. [8851276]. <https://doi.org/10.1109/TMM.2019.2944214>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

The Importance of Context When Recommending TV Content: Dataset and Algorithms

Miklas S. Kristoffersen, Sven E. Shepstone, and Zheng-Hua Tan

Abstract—Home entertainment systems feature in a variety of usage scenarios with one or more simultaneous users, for whom the complexity of choosing media to consume has increased rapidly over the last decade. Users’ decision processes are complex and highly influenced by contextual settings, but data supporting the development and evaluation of context-aware recommender systems are scarce. In this paper we present a dataset of self-reported TV consumption enriched with contextual information of viewing situations. We show how choice of genre associates with, among others, the number of present users and users’ attention levels. Furthermore, we evaluate the performance of predicting chosen genres given different configurations of contextual information, and compare the results to contextless predictions. The results suggest that including contextual features in the prediction cause notable improvements, and both temporal and social context show significant contributions.

I. INTRODUCTION

The underlying factors affecting users’ choices of what to watch on TV have for several years been of interest to commercial and academic research. In the midst of a rapidly changing device and multimedia landscape, TVs continue to be at the core of multimedia consumption in the home with scenarios covering, among others, social gatherings and solitary immersive moments. The inherent complexity of viewing situations challenges the creation of experiences that match personal preferences as well as preferences introduced by the situations themselves. That is, priorities stemming from aspects such as the temporal and social settings of a viewing situation, generally referred to as the context.

Due to the increased availability of multimedia, the focus of recent research has been on improving the users’ decision process by reducing large catalogs of content to a few personalized suggestions [1]. Commercial recommender solutions are now considered core to the business of engaging users and thereby preventing abandonment [2]. To do so, recommender systems have explored various features for personalization, such as history of watching, ratings, user/item similarity, and time of the day, the last of which is an example of features characteristic to context-aware recommender systems (CARS)

M.S. Kristoffersen is with the Research Department of Bang & Olufsen A/S, 7600 Struer, Denmark, and the Signal and Information Processing Section, Department of Electronic Systems, Aalborg University, 9220 Aalborg Ø, Denmark (email: mko@bang-olufsen.dk).

S.E. Shepstone is with the Research Department of Bang & Olufsen A/S, 7600 Struer, Denmark (e-mail: ssh@bang-olufsen.dk).

Z-H. Tan is with the Signal and Information Processing Section, Department of Electronic Systems, Aalborg University, 9220 Aalborg Ø, Denmark (e-mail: zt@es.aau.dk).

The work is supported by Bang & Olufsen A/S, Denmark, and the Innovation Fund Denmark (IFD) under File No. 5189-00009B.

[3]. The main objective of traditional recommender systems is to personalize the experience to the individual, often by studying the user-item matrix. This is an issue, since a TV is often shared by multiple members of a household that engage in both solitary and social viewing situations during everyday life. A solution for accommodating different solitary users on one TV is to allow multiple accounts, but this does not solve the problem when a group of users want a shared experience [2], which is the focus within the research field of group recommender systems [4]. Hence, to achieve recommendations aimed at dynamic compositions of users, the system needs to be aware of the social context it is used in, and thus more features than those classically used for personal predictions are needed [5], [6]. These features can be in the form of personal preferences for each member of the group as seen in [4], e.g. based on social attributes aggregated from social media [7], [8], or a more general social context as seen in [9], e.g. a group of adults watching TV. In this work, we adopt the latter approach and describe social situations with e.g. the number of viewers and their relationship. Recent studies even suggest to decouple the goal of tailoring the experience to the individual (personalize) from tailoring to the situation and intent (contextualize) [10], thereby focusing more on the immediate context than the past behavior of a user. In this study, we investigate some of the advantages and disadvantages of contextless personalization as well as contextualized suggestions. We focus less on algorithmic improvements within each approach, e.g. optimization of contextless personalization, and more on the contribution of different contextual dimensions.

Even though the concept of context-aware recommendations has been studied in several academic and commercial projects, there is still a need for publicly available datasets since only a limited number of such datasets exist, e.g. [11]. Furthermore, the majority of existing CARS datasets are based on explicit feedback, often in the form of ratings. However, within TV recommender systems the feedback is usually implicit in the form of watched/not watched, since continuously probing users for explicit feedback would significantly alter the viewing experience. The use of implicit feedback also means that a user can provide feedback for the same content multiple times, which is typically not the case for e.g. movie ratings.

Another challenge of using existing datasets for developing CARS within the TV domain, is that the most dominant contextual feature is timestamps, i.e. temporal information. Though TV viewing is highly driven by habits linked to temporal context such as time of day, day of week, and season, CARS based exclusively on temporal information

could miss non-trivial correlations between e.g. temporal and social contexts. It is, however, challenging to collect TV consumption data that include contextual information beyond timestamps. People meters¹, for instance, are challenged, [12], by non-compliance (participants neglect to push a button), and secondly, since meters can only log what was shown, and not necessarily whether the participant was engaged, there is no information of participants’ actual exposure to content. In other words, the TV could be showing some content that the participant does not watch.

In this study, we collect and analyze self-reported interactions for a limited number of content classes (genres). We include several contextual features, which allow inspection of patterns in the consumption of different genres. In addition to the comprehensively studied temporal information, our contribution includes novel investigations of associations between consumed content and social settings, e.g. who is present. Using the Experience-Sampling Method (ESM) [13], we ask participants to report TV consumption multiple times each day for a five week period. Through self-reported data, we decrease uncertainty of exposure to content, and allow collection of non-trivial information, such as how much attention is paid to the TV. The data is structured to accommodate quantitative analyses, e.g. in the CARS community, and is publicly available under the name Contextual TV (CTV) dataset². Note that the self-reported information provided by participants could potentially be collected implicitly at run-time in a real-world implementation using technologies such as user detection [14], identification [15], attention level estimation [16], and indoor localization [17].

Using different feature configurations, we also show how well-established methods perform in predicting consumed content given contextual settings, and compare this with context-less prediction. In an initial study, we showed the effectiveness of including contextual information [18], which we expand in the present contribution with in-depth data analysis and detailed investigation of prediction performance. That is, we assess gains of adding contextual knowledge to the prediction task, and study contributions from each contextual dimension to the overall ability to predict which genre a user will engage with in a given situation.

The rest of the paper is organized as follows. We start by surveying related work in Section II. In Section III we introduce and analyze the contextual TV dataset. Section IV presents methods for predicting content and evaluates different configurations. Finally, Section V and VI discuss the findings and conclude the study.

II. RELATED WORK

A. Contextual Aspects of Watching TV

Previous studies of users’ TV watching behavior in given contexts have shown that the TV is mostly a social platform and consumption takes place in a wide variety of scenarios. In

¹A device used to log users’ TV viewing behavior. For identification, the device often relies on participants to push a button on a remote when they enter and leave.

²Available at <http://kom.aau.dk/~zt/online/ContextualTVDataset>.

TABLE I
RELATED CONTEXT-AWARE TV CONTENT RECOMMENDER STUDIES

Ref.	Temp.	Social	Mood/ Emotion	Att.	Dataset pub. avail.
[32]–[35]	✓	–	–	–	–
[11]	✓	–	–	–	✓
[36]	–	–	✓	–	–
[37]	✓	–	✓	–	–
[9], [38], [39]	✓	✓	–	–	–
CTV (ours)	✓	✓	–	✓	✓

[19], 30 households were scan-sampled every 10 minutes for four days, to reveal patterns of who was watching, when, and with whom. Noticeably it was found that 64% of the time, family members were watching TV together. [20] presents the results of an online survey with 550 valid responses. Their results confirm that contextual settings, and not only the users’ personal profiles, are of importance to the decision of what content to watch. In [21] a multi-method field study is conducted with 11 participants. Both temporal and social settings are highlighted as key contextual indicators of consumed content. An example of collecting qualitative contextual TV consumption data using diaries is presented in [22] for 12 households over a three week period. The main difference between the studies listed above and our work, is that we aim for a quantitative dataset that avoids the recall bias associated with questionnaires. To this end, ESM has proven useful for obtaining frequency and patterning of daily activities and social interactions [23]. In a recent study, [24] combined automated data logging from the TV with event-triggered ESM to show, among others, how social context affects TV volume. That is, they used a number of sensors to automatically extract contextual settings of TV viewing events, e.g. Bluetooth trackers to identify present users and their activity level, together with chatbot sessions for obtaining self-reported information of e.g. social context.

B. Recommendations Based on Context

The task of recommending content to users based on their past behavior as well as context, is an active research field. Early work focused on pre- and post-filtering [3], while recent studies have included contextual information directly in the model. The main approaches are tensor factorization [25], factorization machines [26], [27], and most recently efforts based on deep learning such as [28]. In [29] an Android application is used to collect smartphone sensor-based contextual point of interest data from 90 students for a month long period, and recommendations are based on deep auto-encoding. As in this work, they use ESM for collecting feedback from participants. [30] proposes a large-scale context-aware video retrieval method, and [31] presents a method for recommending movies based on emotional preferences of users.

Within TV content recommendation, several studies have based recommendations on context. Table I summarizes related works, and shows which contextual features are included in each study. The table also lists the availability of data from each contribution (to the best of the authors’ knowl-

edge). An early example that includes contextual information is presented in [32] that collected people meter data from approximately 60 participants for one year. Their system has three main components: explicit user modeling based on explicit feedback from users; stereotypical user modeling based on age, gender, etc.; dynamic user modeling based on implicit information inferred from users’ viewing behavior in certain temporal contexts. Recently, [33] and [35] similarly studied temporal aspects of recommending TV content. In [37] approximately 100 participants provided diary data for two weeks. Additional to temporal information, the collected viewing contexts include three mood selections (happy, bored, and unhappy). Feedforward neural networks are used for recommending content, and their results suggest that users’ emotional state helped improve the performance. [38] presents recommendations using support vector machines (SVM) for people meter data collected from 20 families in Finland during a five month period. The viewing contexts consist of temporal and social (additional viewers) information. Their results suggest that including social context makes a minor improvement, but that the improvement depends on family habits, i.e. the correlation between temporal and social settings in families’ typical viewing behavior. [39] also includes social context, using RFID tags to identify users. They evaluate their system in a real-world implementation. In [36], users’ moods are used to improve navigation of programs available in the electronic program guide (EPG). [11] presents a people meter dataset containing implicit viewing events with timestamps for a four-month duration. They smooth temporal context and use distance between contextual settings to recommend TV programs. [9] presents a comparable dataset, but includes *familiar context*, that is, the additional users watching. Their results suggest that temporal context cancels the effect of social context when using both to recommend TV content. In [34] more than 700 million views collected on a Tunisian TV platform are used to recommend TV content in a given context. Their viewing contexts are defined by location, time/day, weather and occasion.

As evident from the literature listed above and in Table I, context-aware TV recommendations have primarily revolved around quantitative data collected through (people) meters. Though meters have a lot of advantages, such as enabling easy large-scale implicit feedback collection, they do (as previously stated) suffer from e.g. non-compliance and actual exposure uncertainty. To better embrace the complexity of viewing situations, we ask participants to provide information that is not easily accessed through meters, such as how much attention a user is paying in a given viewing situation. Also, instead of relying on participants to continuously register their presence in front of the TV using a remote control, our adopted data collection method is chosen to reduce noisy measurements of social settings. Another observation from the literature is that studies have difficulties comparing their findings to those of other works, partially because there is no tradition of sharing results on common datasets, as is the tradition within e.g. movie recommendation. Thus, to the best of the authors’ knowledge, the present contribution is the first to publicly share a dataset with TV viewing events that include

TABLE II
QUESTIONS AND SELECTION OPTIONS IN THE DATASET

Questions	Options
Q1: Have you watched TV within the last four hours?	Yes, no
Q2: Who were you watching it with?	<i>Multiple-option:</i> Alone, partner, child (0-12), child (12+), sibling, parent, friend, other (text)
Q3: How many people (including yourself) watched TV?	1, 2, 3, 4, 5+
Q4: What did you watch?	<i>Multiple-option:</i> News, sport, movie, series, music, documentary, entertainment, children’s, user-generated, other (text)
Q5: Which service(s) did you use?	<i>Multiple-option:</i> Traditional TV, DRTV, TV2 Play, Viaplay, Netflix, HBO Nordic, YouTube, other (text)
Q6: How much attention did you pay to the TV?	None-full (5 steps)

contextual settings beyond timestamps.

According to [40], users tend to choose content such that a few dominant items or providers account for the majority in consumption. They refer to this phenomenon as contextual bias, and discuss how this decreases the diversity of recommenders. In this paper, we include scores that enable assessment of diversity in predictions.

III. CONTEXTUAL TV WATCHING DATASET

This section details the procedure for collecting the dataset, and highlights a number of patterns within the data. The quantitative analysis is focused on general contextual tendencies of viewing situations as well as considerations of temporal and social context dynamics.

A. Experimental Protocol

To obtain data from participants we developed a web page, thereby allowing access from all devices equipped with Internet access and a web browser, though we recommend the use of mobile devices. Participants were asked to answer questions five times every day at 8, 12, 17, 20, and 22 (or when going to bed) for a period of 36 days. These intervals were chosen to accommodate work and study schedules, while still providing ample opportunity to participate over a full day period. Participants were allowed to answer more frequently and at other times than the five pre-specified intervals. Scheduled reporting was preferred to event-based reporting, e.g. asking participants to answer after every TV watching session, since this requires less from the participants to remember and enables evaluation of compliance. Also, scheduled reporting allows signaling to participants to help remind them about the study. Specifically, we used a public calendar with alerts for iOS devices and web push notifications for all other types of devices. Prior to launch, a three-day pilot test was conducted involving 12 potential participants evaluating web page and reminders. Participants for the main study were recruited through social

media with a lottery of three loudspeakers worth €170 each as incentive at the end of the study. A requirement for joining the lottery was at least 14 days of active participation.

The first time a user visits the web page that person is instructed to answer background information questions as part of the enrollment procedure. The collected information includes: Gender, age group, language (Danish/English), device type, household size, additional household members, frequency of TV watching, and favorite TV genres.

On subsequent logins, participants were asked the questions listed in Table II. The questions are designed to have a low cognitive load and take less than 30 seconds to answer. The general flow is that Q2-Q6 are asked only if the selection for Q1 is *yes*. Also, Q3 is skipped if *alone* is selected for Q2. Note that in most cases it is not possible to infer the value of Q3 from Q2. As an example, Q2 option *friend* can be observed with all values for Q3 (except 1) if more friends are present. For Q5 all except *Traditional TV* (and possibly *other*) are streaming services, some specific to Denmark/Scandinavia. The multiple-option questions allow more than one selection, e.g. *partner* and *friend*. Participants are instructed to split answers with different contextual settings, e.g. watching news alone and children’s TV with a child. Answers are logged with the following format: Answer ID, User ID, timestamp, Q1, Q2, Q3, Q4, Q5, Q6. In this study we extract two pieces of information from the timestamp. One is the day of the week that can also be used to determine whether or not it is weekend. Also, we group public holidays and weekend unless stated otherwise. The other feature is the time of the day. We use five groups: 1) Morning: 6-10; 2) Noon: 10-14; 3) Afternoon: 14-18; 4) Evening: 18-22; 5) Night: 22-06.

B. Data Analysis

1) *Participants’ Background*: A total of 118 participants (64 male and 54 female) in the age range 13-70 took part in the ESM study. 57% of the participants were in the age range 21-30, and 84% lived in households with at least two members including themselves. At the enrollment, 81% reported that they watched TV daily and 97% that they watched at least once a week. Concerning reported favorite genres, the *series* genre attracts the largest audience (89). *Movie* (75) and *documentary* (73) are also popular, while *entertainment* (63) and *news* (61) follow as fourth and fifth, respectively. *Sport* (44) is sixth, *user-generated* (21) seventh, and *music* (17) and *children’s* (17) share the last place.

2) *Activity*: Fig. 1 shows the development in enrolled and active (at least one answer that day) participants each day of the 36 days in the study. Notice the relatively large drop-off between enrolled and active participants within the first five days of the study, mainly caused by one-time visitors (a total of 31 throughout the study). From day seven onwards, the number of active participants decreased on average by three every fourth day. The average number of active participants per day was 53, and a total of 60 participants met the requirement of at least 14 active days. Saturdays (day 5, 12, 19, etc. of Fig. 1) had a tendency of fewer active participants, and also showed to be the day with the least responses (935) compared

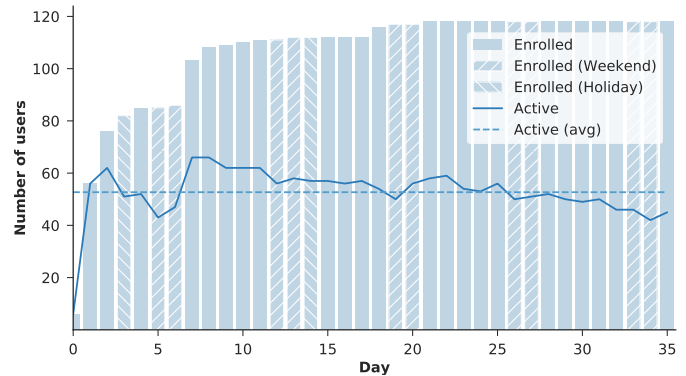


Fig. 1. Number of enrolled and active participants for each day of the study. The average number of active participants per day is shown with a dashed line.

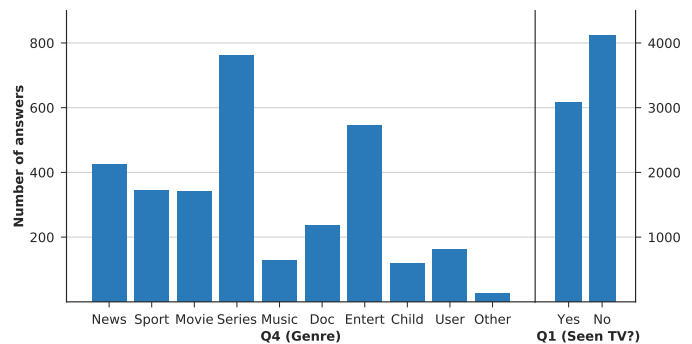


Fig. 2. Count plots for answers to Q4 (left) and Q1 (right).

to the most on Wednesdays (1115). Notably, it seems that part of the inactive users returned Sundays after a “day off”. In terms of time of day, participants answered most frequently in the evening (1784), the least in the morning (1203), and the remaining ranked as follows: night (1468), noon (1418), and afternoon (1328).

3) *TV Consumption*: The dataset consists of 6443 answers. Each answer including more than one selected option for Q4 is split (with the same values for the other entries), which brings the total number of answers to 7201. From these, 3090 are answers with *yes* for Q1. Fig. 2 shows the distribution of answers for Q1 and Q4. It is worth noticing that *series* as the most frequent selection for Q4 accounts for approximately 25% of the answers. This is in accordance with the reported TV favorites in the background information of the participants, where 75% of the participants reported *series* as a favorite. When comparing the reported TV content favorites with Fig. 2, two genres stand out in particular, namely *movie* and *documentary*. These are close competitors for second place among reported TV favorites, but in answers for Q4 they have relatively low counts. This may not come as a surprise, since they typically have a longer duration and might require more attention than the other genres, and thus may not be consumed as frequently. In addition to being the most watched genre, on average participants report to be slightly more attentive when watching *series* compared to both *movie* and *documentary*, as shown in the top of Fig. 3. The reason for the low counts of *movie* and *documentary* needs more analysis, and it would be




Fig. 3. Top: Average attention level (Q6) reported for each genre (Q4). Bottom: Average number of viewers (Q3) reported for each genre (Q4). The black bars indicate \pm one standard deviation.

interesting to study whether it is because users struggle to find and select content within these two genres in particular.

Another point to highlight in Fig. 3 is that *music* and *children's* have low average attention levels compared to the other genres. In the case of *children's* it is most likely because the TV is used primarily by children of the respondents. Also, as seen from the bottom of Fig. 3, *children's* shows mainly to be a social genre. Notice the relatively large standard deviation within the genres *sport* and *music* indicating that these are consumed in different social settings, sometimes by one user and at other times by groups of users.

Pearson's chi-square test is used to measure the level of association between the choice of genre (removing answers with the selection *other*) and the contextual features. Also, Cramér's V is reported. To this end, a contingency table is formed for each contextual dimension, and results are presented in Table III for all cases where at least 80% of the expected frequencies are above five and none are zero. Statistical significant interactions are found for all features.

TABLE III
PEARSON'S CHI-SQUARE TESTS OF ASSOCIATION BETWEEN CHOICE OF GENRE AND CONTEXTUAL FEATURES

Feature	χ^2	dof	p	V
Time of day	326.39	32	< 0.001	0.16
Weekday/[weekend/holiday]	124.52	8	< 0.001	0.20
Additional viewers	1192.53	56	< 0.001	0.21
Number of viewers	540.47	32	< 0.001	0.21
Attention level	593.36	32	< 0.001	0.22
Service	2169.12	56	< 0.001	0.29

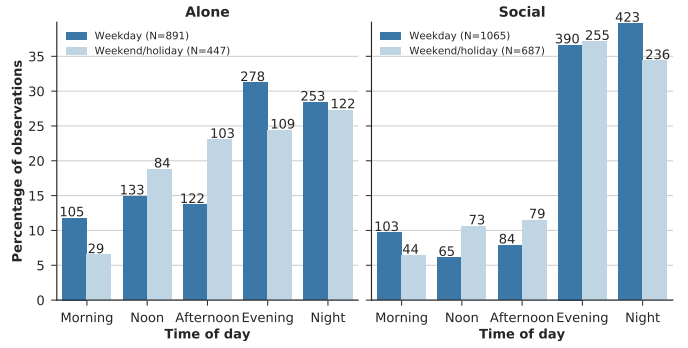


Fig. 4. The distribution of TV consumption according to time of day in four contextual settings, alone/social and weekday/weekend, such that e.g. alone+weekday sums to 100%. The numbers above the bars indicate the actual number of observations.

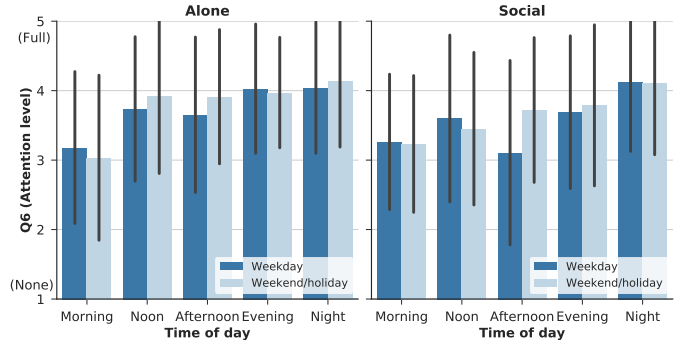


Fig. 5. The average attention level according to time of day in four contextual settings, alone/social and weekday/weekend. The black bars indicate \pm one standard deviation.

4) *Temporal and Social Aspects of TV Consumption:* We have shown that the choice of genre is associated with several contextual settings. However, questions, such as what changes during the day, in the weekend, or in social situations, remain unanswered, and are hence investigated below.

A fundamental difference is the consumption pattern. As shown in Fig. 4, TV watching in social settings (with at least one co-viewer) happens most frequently during the evening or night both for workdays (75%) and weekends (70%). Morning, noon, and afternoon account for the remaining 25% of observations for weekdays and 30% for weekends in social situations. TV watching in solitary settings is more spread throughout the day, though weekdays show the same tendency with evening and night being the dominant time slots. Independent of social context, TV watching in the morning occurs most frequently during weekdays, while the share of noon and afternoon viewing increases in weekends. Also note from the figure that approximately 57% of all observations (3090) take place in a social context (1752), while 63% are during workdays (1956).

Another element to consider is the attention level of the users as presented in Fig. 5. Generally, users pay more attention to the content as the day progresses. A notable exception is social afternoons in weekdays that have the lowest average attention level among all social settings, which could possibly be because users have just returned back home from work and engage in conversations and other activities while watching




Fig. 6. The distribution of genres consumed in four contextual settings, alone/social and weekday/weekend. The numbers above the bars indicate the actual number of observations.

TV. There is also a tendency that users pay more attention when they are alone. Exceptions are mornings and nights, where the levels are approximately similar. Interestingly, when compared to watching alone, it seems that the way users co-view TV changes between afternoon, evening, and night, such that the social activity tends to be more focused around watching TV intensively as it gets late.

Lastly, Fig. 6 shows how the temporal and social context influence the choice of genre. Note that the height of the bars indicates the consumption share of a genre within one of the four contextual settings, alone+weekday, alone+weekend, social+weekday, and social+weekend. Hence, it can for example not be concluded that *movie* is watched more in weekends than weekdays in social settings (actually, the opposite is the case). It can, however, tell that the proportion of times *movie* is selected over other genres is higher in the weekend. A number of genres show clear differences between weekdays and weekends. *News*, *series*, and *entertainment* are preferred during weekdays, while *sport* and *movie* increase their share considerably during weekends. Though *series* are watched less in the weekends, it is the genre with the largest share for all contexts. *Entertainment* is second in three out of four contexts, but drops behind *sport* and *movie* to fourth for social weekends. *News* follows just behind *entertainment* in weekdays, but it is the genre that decreases most in weekends. *Movie* and *children’s* are preferred in social settings, while *user-generated* is mainly consumed when alone, which could possibly be because the genre has its main roots on smaller screens, where users mainly consume it in solitary settings. The proportion for *music* is similar among the four contexts.

IV. PREDICTION OF PREFERENCES

In this section we present how the contextual features of the CTV dataset can be used for prediction of consumed content. The rationale behind is that not only will a user’s viewing history influence the choice, but so will the circumstances in which that user is watching TV, which is motivated by the findings of the previous section that shows how contextual settings influence viewing situations contained in the collected dataset. The procedure and experiments presented in this

section provide insights into the mapping between multimedia content genres available on modern TVs and users’ personal preferences, and how this mapping can be coupled with context information.

The goal of the methods described in this section is to predict what genre a user is going to watch in the reported context. Specifically, let \mathcal{Y} be the set of possible genres to choose from in Q4 (see Table II). Then \hat{y} is the predicted genre using a trained classifier, f , with input $\mathbf{x} \in \mathbb{R}^D$ consisting of user and context features:

$$\hat{y} = f(\mathbf{x}), \text{ where } f: \mathbb{R}^D \rightarrow \mathcal{Y}. \quad (1)$$

Various methods can be applied to this challenge, such as state-of-the-art factorization machines [27] or neural networks [28]. The focus of this investigation is on the contribution of the contextual dimensions and the type of errors, for which reason we do not present the results of a wide palette of algorithms, but rely on a few that are well-known within the machine learning community for establishing baselines of the CTV dataset.³ We do, however, show how temporal and social context affect the prediction ability, and compare it to e.g. contextless prediction.

A. Features and Methods

The task, as shown in (1), is defined as a multiclass classification problem with the users’ selections for Q4 as target. The selections for the remaining questions are used as contextual features (see Table IV). All features are categorical and represented using one-hot encoding. Neither user demographics, nor the optional text input for *other* in Q2, Q4, and Q5 are included in this study.

Table IV also lists a number of feature configurations. In this work, we define the feature configuration *all* as service-independent. Thus, *all* is a collection of all features except the *service* feature. The reasoning behind this definition is that a prediction or recommendation of genre will have most impact across providers, since some services have a very targeted range of content genres, e.g. YouTube relies heavily

³The dataset includes code for the baselines presented in this section.

TABLE IV
FEATURES AND EXAMPLES OF FEATURE CONFIGURATIONS

Feature	Dimensions	Origin
(U) User ID	118	Login
(T) Time of day	5	Timestamp
(D) Weekday/weekend	2	Timestamp
(W) Additional viewers	8	Q2
(M) Number of viewers	5	Q3
(S) Service	8	Q5
(A) Attention level	5	Q6
Feature Configurations		
<i>all</i>	143	UTDWMA
<i>all+S</i>	151	UTDWMSA
<i>all-U</i>	25	TDWMA

on the user-generated genre. The *service* feature is included in the *all+S* configuration. Other configurations can be used as well, e.g. time of day and weekday/weekend (*TD*). A notable configuration is what we refer to as the contextless, which consists of purely user identity information (*U*).

We include six approaches in the experiment to establish baselines for the CTV dataset. The methods are selected based on their well-documented achievements in numerous domains. Four methods are compared using scikit-learn [41] implementations: logistic regression (LR), gradient boosting decision trees (GBDT), support vector machines (SVM), multi-layer perceptrons (MLP). Additionally, we include the most popular (toppop) and random (random) predictors. For toppop, genres are ranked by their popularity judged by the number of observations in the training set. The random predictor randomly ranks the genres for each prediction.

B. Evaluation

The methods are evaluated using nested cross-validation (also referred to as double cross-validation) with five outer folds and three inner folds. That is, in the outer loop the dataset is split into five folds, using one fold in turn as test set and the remaining four folds as training set. The training set for each outer iteration is further divided into three inner folds for optimization of hyperparameters. To this end, on a rotational basis, two (inner) folds are used for training and one is used as validation set. The best scoring hyperparameter configuration of the inner loop is used to assess the predictive performance of the model on the outer test fold by training on the full training set. We report the average performance across the outer folds and the standard deviation. Users that had not answered at least five times were not included in the evaluation.

1) *Configuration of Hyperparameters*: Due to considerations of computational complexity, some hyperparameters are determined empirically prior to run-time and are static throughout feature configurations. We fit the LR weights using stochastic average gradient descent with L2 regularization, and set the multi-class parameter to "multinomial" for softmax regression. For GBDT, we use 1000 boosting stages, each fitted on a random subsample consisting of 50% of the training samples. The SVM use a one-vs-rest decision scheme, and

MLP is implemented with two hidden layers each consisting of 200 neurons with rectified linear (ReLU) activation functions and a softmax output layer, optimized during training using Adam and L2 regularization of weights. During hyperparameter tuning in the inner loop of the nested cross-validation, the following variables are determined: LR - regularization strength; GBDT - maximum depth of individual trees; SVM - kernel type (linear/RBF), kernel coefficient and regularization strength; MLP - regularization strength.

2) *Metrics*: The hit ratio at *K* predictions (HR@*K*) is used as a metric for evaluation. At *K* larger than one, multiple guesses are allowed for each trial. It is calculated using:

$$\text{HR@K} = \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^K \mathbf{1}(\hat{y}_{n,k} = y_n) \quad (2)$$

where *N* is the number of trials. $\mathbf{1}$ is the hit/indicator function, which is one if the prediction, $\hat{y}_{n,k}$, is equal to the actual target, y_n , and zero otherwise. The \hat{y}_n is sorted with predictions in ascending order according to confidence score, such that $\hat{y}_{n,1}$ is the most probable prediction. The mean reciprocal rank (MRR) is used to assess the average ranking (*k*) of the true targets:

$$\text{MRR} = \frac{1}{N} \sum_{n=1}^N \sum_{k=1}^M \frac{1}{k} \mathbf{1}(\hat{y}_{n,k} = y_n) \quad (3)$$

where *M* is the total number of target classes. We also report F1 scores with macro averaging. F1 (macro) is an average of each individual class performance. It gives equal weight to classes, which means it can be used to assess the performance on small classes. This makes F1 (macro) an indicator of a method's ability to predict diverse target classes. On the other hand, F1 (micro) pools all trials with equal weight. Thus, classes with many samples will dominate classes with few samples. F1 (micro) is with this setup (multi-class single-label) equal to HR@1 and therefore not presented explicitly.

C. Results

1) *Feature configuration all*: The performance using feature configuration *all* is shown in Table V. It is not a surprise that random scores approximately 0.1 in HR@1 (10 target classes), and that toppop scores close to 0.25 since it was shown that *series* as the most watched genre accounts for approximately one fourth of the data points. Also, toppop outperforms random in terms of HR@3 and MRR. Note, however, that random performs better than toppop for F1 (macro), due to the diversity in predicted genres.

The remaining methods achieve considerably higher scores than random and toppop. LR, as the best scoring, almost doubles the HR@1 of toppop, and successfully predicts the genre in approximately 44% of the cases and 82% when allowing three guesses.

2) *Genre confusions*: A deeper look into the prediction errors of LR for feature configuration *all* is shown in a confusion matrix in Fig. 7. *Series* receives many predictions (1022) compared to the actual number of occurrences (740) resulting in the best recall score of all the genres (true positives over number of observations, $528/740 = 0.71$). Despite the

	Predicted										Sum	Recall
	News	Sport	Movie	Series	Music	Doc	Entert	Child	User	Other		
News	158	52	22	49	7	16	91	14	2	0	411	0.38
Sport	65	117	27	42	7	2	56	12	9	0	337	0.35
Movie	26	29	76	131	3	5	42	19	7	0	338	0.22
Series	25	22	39	528	5	12	88	3	17	1	740	0.71
Music	8	11	6	17	64	0	10	0	10	1	127	0.50
Doc	51	8	14	70	4	15	62	3	1	0	228	0.07
Entert	83	38	34	128	7	5	222	14	10	0	541	0.41
Child	5	5	9	1	1	1	7	87	1	0	117	0.74
User	8	7	7	45	23	0	12	3	54	0	159	0.34
Other	1	1	3	11	2	0	2	4	0	1	25	0.04
Sum	430	290	237	1022	123	56	592	159	111	3	3023	
Prec	0.37	0.40	0.32	0.52	0.52	0.27	0.38	0.55	0.49	0.33		
F1	0.38	0.37	0.26	0.60	0.51	0.11	0.39	0.63	0.40	0.07		

Fig. 7. Confusion matrix using LR for feature configuration *all* and the aggregated (out-of-fold) cross-validation predictions. The bottom rows show the precision and F1 for each genre, and the right-most column shows the recall.

many predictions, it also manages to achieve one of the highest precision scores (true positives over number of predictions, $528/1022 = 0.52$). The opposite is the case with *documentary*, which receives few predictions (56) compared to observations (228), leading to a low recall score (0.07). Together with the precision score of 0.27, the resulting F1 score is 0.11, which emphasizes that *documentary* is difficult to predict in this specific setup. It can be seen from the off-diagonal entries that there are three main genres that are confused with *documentary*, being *news*, *series*, and *entertainment*.

Two genres stand out in terms of F1 score, namely *series* and *children’s* that both achieve a score around 0.6. Next is *music* with a score of 0.5, which gets a majority of its false positives when the observed genre is *user-generated*. The *user-generated* genre has a high (relative) precision, but a low recall causes the F1 score to fall to an average level among the genres, meaning that even though it is hard to retrieve *user-generated* content for recommendation, when it is finally predicted, the predictions are fairly reliable. *News*, *sport*, and *entertainment* also have F1 scores that are close to the F1 (macro). In addition to *documentary*, the *movie* genre underperforms. The F1 score of *movie* is low mainly due to a low recall, caused to a high degree by *series*, but it is also frequently confused with *news*, *sport*, and *entertainment*.

3) *Contextual dimensions*: A comparison of the methods for multiple feature configurations is shown in Fig. 8. Note that the random and toppop scores are independent of feature selection. Also, as can be seen from the figure, in general

TABLE V
RESULTS FOR THE FEATURE CONFIGURATION *all*
(STANDARD DEVIATION IN PARENTHESES)

Method	HR@1	HR@3	F1 (macro)	MRR
random	0.101 (0.005)	0.295 (0.004)	0.087 (0.008)	0.290 (0.005)
toppop	0.245 (0.009)	0.560 (0.014)	0.039 (0.001)	0.460 (0.008)
MLP	0.413 (0.018)	0.787 (0.013)	0.335 (0.022)	0.621 (0.012)
GBDT	0.417 (0.014)	0.786 (0.021)	0.354 (0.027)	0.623 (0.009)
SVM	0.425 (0.009)	0.809 (0.021)	0.358 (0.015)	0.632 (0.012)
LR	0.437 (0.026)	0.815 (0.010)	0.373 (0.031)	0.641 (0.013)

TABLE VI
RELATIVE PERFORMANCE WHEN EACH CONTEXTUAL VARIABLE IS REMOVED COMPARED TO FEATURE CONFIGURATION *all* USING LR

Feat. Conf.	all-D	all-M	all-T	all-A	all-W	all-U
Rel. HR@1	-0.9%	-1.4%	-2.1%	-2.3%	-2.5%	-29.5%
Rel. F1 (macro)	-2.4%	-3.8%	-5.9%	-5.1%	-4.0%	-44.2%

there is not a large deviation between performance of the four top performing methods within each feature configuration. Therefore, we mainly highlight results between feature configurations in the following.

The worst performing feature configuration in Fig. 8 is the temporal configuration, *TD*, for which the performance of LR almost decreases to comparable results of random and toppop, though the HR@1, HR@3, and MRR are significant better than those of the random predictor, and likewise the F1 (macro) score is higher than what toppop achieves. Adding social aspects (*TD*→*TDW*) significantly improves performance according to McNemar’s test⁴ (for LR: $\chi^2(1)=29.10$, $p<0.001$, $V=0.09$). Furthermore, knowledge of additional viewers and attention level, *WA*, scores on par with *TDW*, and even slightly improves in terms of F1 (macro).

Notably, contextless prediction, *U*, outperforms *all-U*, which indicates the importance of user specific behavior and habits (for LR: $\chi^2(1)=26.33$, $p<0.001$, $V=0.09$). The F1 (macro), however, does not see as large an improvement between the two configurations as the three other metrics. Adding temporal context to the user ID (*U*→*UTD*) increases performance slightly, while F1 (macro) improves considerably suggesting that the temporal information enhances the ability to predict diverse genres. As is the case with *TD*→*WA*, *UWA* performs significantly better than *UTD* (for LR: $\chi^2(1)=9.83$, $p\approx 0.002$, $V=0.06$). In fact, *UWA* provides results on the same level as *all* with no significant difference according to McNemar’s test (for LR: $\chi^2(1)=3.52$, $p>0.05$, $V=0.03$). Lastly, knowing which service a user is using, *all+S*, achieves the highest scores.

The contribution of contextual dimensions to the overall result is investigated in Table VI, which shows the average performance when each variable is removed from the LR model in turn. The results are relative to LR for the feature configuration *all*. Again, user identity (*U*) shows to constitute a vital component of the model with substantial performance reductions compared to the other features. On the contrary, the weekday/weekend feature (*D*) contributes the least among the inputs. Interestingly, the metrics are affected differently for the remaining features. Removing information of additional viewers (*W*) reduces HR@1 the most, while the time of day (*T*) has the largest impact on F1 (macro). This indicates distinct strengths among the features. For example, the time of day helps predict a diverse range of genres, which is also suggested by the previously discussed *U*→*UTD* results of Fig. 8.

⁴A matrix, $A_{2 \times 2}$, is formed with $a_{1,1}$ being the number of trials where both methods are correct, $a_{1,2}$ and $a_{2,1}$ are the trials where one of the methods fail, and $a_{2,2}$ is when both are incorrect. McNemar’s χ^2 test statistic and Cramér’s V are then computed as:
 $\chi^2 = (a_{1,2} - a_{2,1})^2 / (a_{1,2} + a_{2,1})$, $V = \sqrt{\chi^2 / \sum_{i,j} a_{i,j}}$.

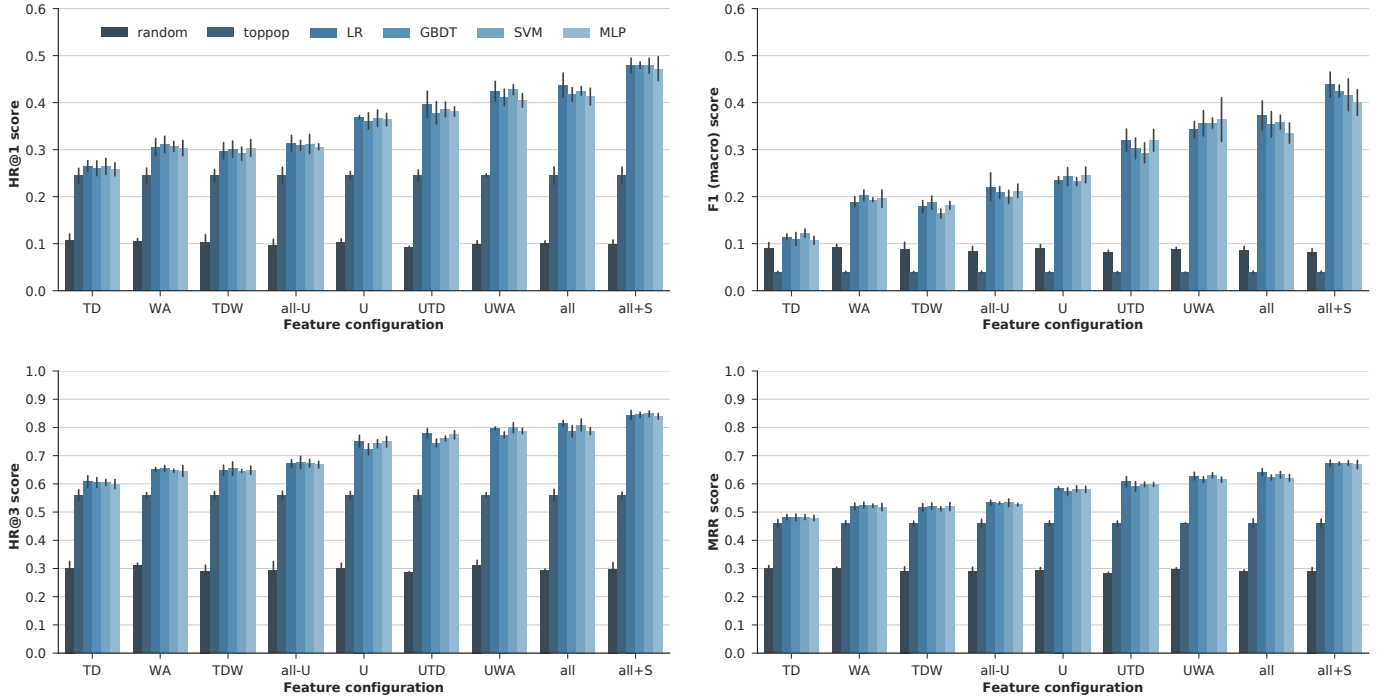


Fig. 8. HR@1, HR@3, F1 (macro), and MRR scores for the six methods using different feature configurations (see Table IV). The height of the bars indicate the average of the cross-validation results and the black bars show \pm one standard deviation.

V. DISCUSSION

The predictive performance using LR for feature configuration *all* (see Fig. 7) is best for the two genres *series* and *children's* in terms of F1 score. According to the data analysis (presented in Section III-B), *children's* exhibit distinctive contextual preferences, such as low attention level and multiple viewers, resulting in fewer prediction confusions. The opposite is the case with *documentary* and *movie* that are not easily distinguishable from the features collected in this study, which also makes them the genres with the lowest F1 scores. A reason for the worse performance when compared to *series*, could be the imbalance of the dataset, and the overall score of the less viewed genres could possibly improve by handling this. Also, one should keep genre ambiguities in mind when evaluating the system. For example, a movie could be a children's movie, thereby potentially covering two genres in this study. Assuming that a participant provides consistent labels throughout the study, these ambiguities would primarily result in reduced performance when making use of information across users.

The analysis of the collected data suggests that contextual aspects are an integral part of users' decision process when selecting what content to watch on TV. This hypothesis is supported by the findings of Section IV, in which it is shown that the inclusion of contextual information is positively associated with improved predictive performance, both in terms of accuracy and diversity. Note that this is without inducing prior expert knowledge of certain situations, such as the often easily distinguishable *children's* setting. As previously mentioned, user demographics are not included in the experiments, but they can potentially enhance the system further. However, brief

experiments showed that information of age group and gender did not introduce measurable performance improvements of LR when user identity is also included. That being noted, the results of Section IV also highlight the importance of knowing who is watching, since *all-U* \rightarrow *U* shows significant improvement. Interestingly, though adding temporal information to the user ID (*U* \rightarrow *UTD*) only shows slight improvement in accuracy, it clearly indicates more diverse predictions. Hence, contextless prediction (*U*) achieves high accuracy scores, but it seems to be at the expense of the ability to predict diverse genres. However, from the results obtained in this work, contextual prediction benefits from knowing user habits, and drawing on contextual patterns of viewing situations should thus not exclude personalization based on past behavior of each individual.

As opposed to the findings of [9] that the effects of social context are canceled when also considering temporal context, our results suggest a significant improvement when adding social context to temporal, e.g. *TD* \rightarrow *TDW*. This is also evident in terms of diversity in the predicted genres. What users watch, when, and with whom may be correlated, shown by the habitual behavior of users when watching TV, but our results indicate that knowing the social context of a viewing situation will enable the system to adapt to some scenarios deviating from temporal habits. As an example, content chosen for Friday nights could be very different depending on the social company of the user, while Monday mornings frequently consist of the same users and content.

VI. CONCLUDING REMARKS

In this paper, we introduced the novel and publicly available CTV dataset of TV consumption enriched with contextual

information that contributes to the evaluation of TV viewing in the home. To this end, we conducted an extensive field study over a period of five weeks with a group of more than 100 participants. Using the dataset, we showed associations between different aspects of TV watching, e.g. users' average attention level of genres, and how these change over time and in different social situations. Furthermore, we evaluated to which degree contextual knowledge influences the performance of predicting what content will be consumed. The experimental results showed that inclusion of contextual information significantly improves accuracy and diversity compared to contextless predictions, but also that knowledge of past behavior is essential to achieve high accuracies.

In future work we plan to apply methods that have proven successful within the context-aware recommender systems community, and evaluate how they adapt to a dataset with a limited number of target classes and multiple interactions between the same user and target. This will among others include deep models suitable for inferring latent contextual features. The presented work is directly applicable to recommending genres. Future work can potentially investigate using genres for reducing the search space in item recommendations.

REFERENCES

- [1] D. V eras, T. Prota, A. Bispo, R. Prud encio, and C. Ferraz, "A literature review of recommender systems in the television domain," *Expert Systems with Applications*, vol. 42, no. 22, pp. 9046–9076, Dec. 2015.
- [2] C. A. Gomez-Uribe and N. Hunt, "The Netflix Recommender System: Algorithms, Business Value, and Innovation," *ACM Trans. Manage. Inf. Syst.*, vol. 6, no. 4, pp. 13:1–13:19, Dec. 2015.
- [3] G. Adomavicius and A. Tuzhilin, "Context-Aware Recommender Systems," in *Recommender Systems Handbook*. Springer, Boston, MA, 2015, pp. 191–226.
- [4] J. Masthoff, "Group Recommender Systems: Aggregation, Satisfaction and Group Attributes," in *Recommender Systems Handbook*. Springer, Boston, MA, 2015, pp. 743–776.
- [5] Y. Shi, M. Larson, and A. Hanjalic, "Collaborative Filtering Beyond the User-Item Matrix: A Survey of the State of the Art and Future Challenges," *ACM Comput. Surv.*, vol. 47, no. 1, pp. 3:1–3:45, May 2014.
- [6] P. Cui, W. Zhu, T. Chua, and R. Jain, "Social-Sensed Multimedia Computing," *IEEE MultiMedia*, vol. 23, no. 1, pp. 92–96, Jan. 2016.
- [7] P. Cui, Z. Wang, and Z. Su, "What Videos Are Similar with You?: Learning a Common Attributed Representation for Video Recommendation," in *Proceedings of the 22nd ACM International Conference on Multimedia*, ser. MM '14. New York, NY, USA: ACM, 2014, pp. 597–606.
- [8] M. Bertini, A. Del Bimbo, A. Ferracani, F. Gelli, D. Maddaluno, and D. Pezzatini, "A Novel Framework for Collaborative Video Recommendation, Interest Discovery and Friendship Suggestion Based on Semantic Profiling," in *Proceedings of the 21st ACM International Conference on Multimedia*, ser. MM '13. New York, NY, USA: ACM, 2013, pp. 451–452.
- [9] P. Cremonesi, P. Modica, R. Pagano, E. Rabosio, and L. Tanca, "Personalized and Context-Aware TV Program Recommendations Based on Implicit Feedback," in *E-Commerce and Web Technologies*, ser. Lecture Notes in Business Information Processing. Springer, Cham, Sep. 2015, pp. 57–68.
- [10] R. Pagano, P. Cremonesi, M. Larson, B. Hidasi, D. Tikk, A. Karatzoglou, and M. Quadrana, "The Contextual Turn: From Context-Aware to Context-Driven Recommender Systems," in *Proceedings of the 10th ACM Conference on Recommender Systems*, ser. RecSys '16. ACM, 2016, pp. 249–252.
- [11] R. Turrin, A. Condorelli, P. Cremonesi, and R. Pagano, "Time-based TV programs prediction," in *1st Workshop on Recommender Systems for Television and Online Video at ACM RecSys*, 2014, p. 7.
- [12] B. Jardine, J. Romaniuk, J. G. Dawes, and V. Beal, "Retaining the primetime television audience," *European Journal of Marketing*, vol. 50, no. 7/8, pp. 1290–1307, Jul. 2016.
- [13] R. Larson and M. Csikszentmihalyi, "The Experience Sampling Method," *New Directions for Methodology of Social & Behavioral Science*, vol. 15, pp. 41–56, 1983.
- [14] D. T. Nguyen, W. Li, and P. O. Ogunbona, "Human detection from images and videos: A survey," *Pattern Recognition*, vol. 51, pp. 148–175, Mar. 2016.
- [15] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A Unified Embedding for Face Recognition and Clustering," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2015, pp. 815–823.
- [16] A. Coifman, P. Rohoska, M. Kristoffersen, S. Shepstone, and Z.-H. Tan, "Subjective Annotations for Vision-based Attention Level Estimation," in *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, ser. VISAPP '19, vol. 5. SciTePress, 2019, pp. 249–256.
- [17] J. Xiao, Z. Zhou, Y. Yi, and L. M. Ni, "A Survey on Wireless Indoor Localization from the Device Perspective," *ACM Comput. Surv.*, vol. 49, no. 2, pp. 25:1–25:31, Jun. 2016.
- [18] M. S. Kristoffersen, S. E. Shepstone, and Z.-H. Tan, "A Dataset for Inferring Contextual Preferences of Users Watching TV," in *Proceedings of the 26th Conference on User Modeling, Adaptation and Personalization*, ser. UMAP '18. ACM, 2018, pp. 367–368.
- [19] D. Saxbe, A. Graesch, and M. Alvik, "Television as a Social or Solo Activity: Understanding Families' Everyday Television Viewing Patterns," *Communication Research Reports*, vol. 28, no. 2, pp. 180–189, Apr. 2011.
- [20] J. Abreu, P. Almeida, B. Teles, and M. Reis, "Viewer Behaviors and Practices in the (New) Television Environment," in *Proceedings of the 11th European Conference on Interactive TV and Video*, ser. EuroITV '13. ACM, 2013, pp. 5–12.
- [21] K. Mercer, A. May, and V. Mitchel, "Designing for video: Investigating the contextual cues within viewing situations," *Personal and Ubiquitous Computing*, vol. 18, no. 3, pp. 723–735, Mar. 2014.
- [22] J. Vanattenhoven and D. Geerts, "Contextual aspects of typical viewing situations: A new perspective for recommending television and video content," *Personal and Ubiquitous Computing*, vol. 19, no. 5-6, pp. 761–779, Aug. 2015.
- [23] M. Csikszentmihalyi and R. Larson, "Validity and Reliability of the Experience-Sampling Method," in *Flow and the Foundations of Positive Psychology*. Springer, Dordrecht, 2014, pp. 35–54.
- [24] M. Kim, J. Kim, S. Han, and J. Lee, "A Data-driven Approach to Explore Television Viewing in the Household Environment," in *Proceedings of the 2018 ACM International Conference on Interactive Experiences for TV and Online Video*, ser. TVX '18. ACM, 2018, pp. 89–100.
- [25] E. Frolov and I. Oseledets, "Tensor methods and recommender systems," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 7, no. 3, p. e1201, May 2017.
- [26] S. Rendle, "Factorization Machines," in *2010 IEEE International Conference on Data Mining*, Dec. 2010, pp. 995–1000.
- [27] S. Rendle, Z. Gantner, C. Freudenthaler, and L. Schmidt-Thieme, "Fast Context-aware Recommendations with Factorization Machines," in *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval*, ser. SIGIR '11. ACM, 2011, pp. 635–644.
- [28] H.-T. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir, R. Anil, Z. Haque, L. Hong, V. Jain, X. Liu, and H. Shah, "Wide & Deep Learning for Recommender Systems," in *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems*, ser. DLRS 2016. ACM, 2016, pp. 7–10.
- [29] M. Unger, B. Shapira, L. Rokach, and A. Bar, "Inferring Contextual Preferences Using Deep Auto-Encoding," in *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, ser. UMAP '17. ACM, 2017, pp. 221–229.
- [30] Y. Feng, P. Zhou, J. Xu, S. Ji, and D. O. Wu, "Video Big Data Retrieval Over Media Cloud: A Context-aware Online Learning Approach," *IEEE Transactions on Multimedia*, pp. 1–1, 2019.
- [31] S. Benini, L. Canini, and R. Leonardi, "A Connotative Space for Supporting Movie Affective Recommendation," *IEEE Transactions on Multimedia*, vol. 13, no. 6, pp. 1356–1370, Dec. 2011.
- [32] L. Ardissono, C. Gena, P. Torasso, F. Bellifemine, A. Difino, and B. Negro, "User Modeling and Recommendation Techniques for Personalized Electronic Program Guides," in *Personalized Digital Television*, ser. Human-Computer Interaction Series. Springer, Dordrecht, 2004, pp. 3–26.
- [33] M. Aharon, E. Hillel, A. Kagian, R. Lempel, H. Makabee, and R. Nisim, "Watch-It-Next: A Contextual TV Recommendation System," in

Machine Learning and Knowledge Discovery in Databases, ser. Lecture Notes in Computer Science. Springer, Cham, Sep. 2015, pp. 180–195.

- [34] M. Bambia, M. Boughanem, and R. Faiz, “Exploring Current Viewing Context for TV Contents Recommendation,” in *2016 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, Oct. 2016, pp. 272–279.
- [35] Y. Park, J. Oh, and H. Yu, “RecTime: Real-Time recommender system for online broadcasting,” *Information Sciences*, vol. 409–410, pp. 1–16, Oct. 2017.
- [36] S. E. Shepstone, Z.-H. Tan, and S. H. Jensen, “Using Audio-Derived Affective Offset to Enhance TV Recommendation,” *IEEE Transactions on Multimedia*, vol. 16, no. 7, pp. 1999–2010, Nov. 2014.
- [37] S. H. Hsu, M.-H. Wen, H.-C. Lin, C.-C. Lee, and C.-H. Lee, “AIMED- A Personalized TV Recommendation System,” in *Interactive TV: A Shared Experience*, ser. Lecture Notes in Computer Science. Springer, Berlin, Heidelberg, May 2007, pp. 166–174.
- [38] E. Vildjiounaite, V. Kyllönen, T. Hannula, and P. Alahuhta, “Unobtrusive dynamic modelling of TV programme preferences in a Finnish household,” *Multimedia Systems*, vol. 15, no. 3, pp. 143–157, Jul. 2009.
- [39] S. Song, H. Moustafa, and H. Afifi, “Advanced IPTV Services Personalization Through Context-Aware Content Recommendation,” *IEEE Transactions on Multimedia*, vol. 14, no. 6, pp. 1528–1537, Dec. 2012.
- [40] F. Lorenz, J. Yuan, A. Lommatzsch, M. Mu, N. Race, F. Hopfgartner, and S. Albayrak, “Countering Contextual Bias in TV Watching Behavior: Introducing Social Trend As External Contextual Factor in TV Recommenders,” in *Proceedings of the 2017 ACM International Conference on Interactive Experiences for TV and Online Video*, ser. TVX ’17. ACM, 2017, pp. 21–30.
- [41] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, “Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research*, vol. 12, no. Oct, pp. 2825–2830, 2011.



Miklas Strøm Kristoffersen (StM’17) received the B.Sc. and M.Sc. degrees in electronic engineering with specialization in vision, graphics, and interactive systems from Aalborg University, Aalborg, Denmark, in 2014 and 2016, respectively, where he now pursues the Ph.D. degree. He is currently a research fellow with Bang & Olufsen A/S, Struer, Denmark. His research interests include computer vision, machine learning, and context-aware recommender systems.



Sven Ewan Shepstone (M’11) received the B.S. and M.S. degrees in electrical engineering from the University of Cape Town, Cape Town, South Africa, in 1999 and 2002, respectively, and the Ph.D. degree from Aalborg University, Aalborg, Denmark, in 2015.

From 2005 to 2010, he researched in the field of broadband communications with Ericsson a/s, Copenhagen, Denmark, and has been with Bang & Olufsen A/S, Struer, Denmark, since 2010, where he is currently a Research Specialist. His current research interest includes AI applied to consumer electronics. Dr. Shepstone was a recipient of the IEEE Ganesh N. Ramaswamy Memorial Student Grant at ICASSP 2015.



Zheng-Hua Tan (M00SM06) received the B.Sc. and M.Sc. degrees in electrical engineering from Hunan University, Changsha, China, in 1990 and 1996, respectively, and the Ph.D. degree in electronic engineering from Shanghai Jiao Tong University, Shanghai (SJTU), China, in 1999.

He is a Professor in the Department of Electronic Systems and a Co-Head of the Centre for Acoustic Signal Processing Research (CASPR) at Aalborg University, Aalborg, Denmark. He was a Visiting Scientist at the Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology (MIT), Cambridge, USA, an Associate Professor in the Department of Electronic Engineering at SJTU, Shanghai, China, and a postdoctoral fellow in the Department of Computer Science at KAIST, Daejeon, Korea. His research interests include machine learning, deep learning, pattern recognition, speech and speaker recognition, noise-robust speech processing, multimodal signal processing, and social robotics. He has authored/coauthored about 200 publications in refereed journals and conference proceedings. He is a member of the IEEE Signal Processing Society Machine Learning for Signal Processing Technical Committee (MLSP TC). He has served as an Editorial Board Member/Associate Editor for Computer Speech and Language, Digital Signal Processing, and Computers and Electrical Engineering. He was a Lead Guest Editor of the IEEE Journal of Selected Topics in Signal Processing and a Guest Editor of several journals including Neurocomputing. He is the General Chair for IEEE MLSP 2018 and was a Technical Program Co-Chair for IEEE Workshop on Spoken Language Technology (SLT 2016).