

Just browsing?

Understanding user journeys in online TV

Yehia Elkhatib^{*}, Rebecca Killick[◊], Mu Mu^{*}, and Nicholas Race^{*}

^{*}School of Computing and Communications, [◊]Department of Mathematics and Statistics
Lancaster University, UK
{i.lastname}@lancaster.ac.uk

ABSTRACT

Understanding the dynamics of user interactions and the behaviour of users as they browse for content is vital for advancements in content discovery, service personalisation, and recommendation engines which ultimately improve quality of user experience. In this paper, we analyse how more than 1,100 users browse an online TV service over a period of six months. Through the use of model-based clustering, we identify distinctive groups of users with discernible browsing patterns that vary during the course of the day.

Categories and Subject Descriptors

H.5.1 [Multimedia Information Systems]: Evaluation/methodology; H.5.2 [User Interfaces]: Interaction styles

Keywords

Online TV, User behaviour, Modelling

1. INTRODUCTION

Consumers are exhibiting an increasing appetite for accessing TV content over the Internet, either as live or catchup/on-demand (VoD) streaming services. In 2013, over a third (36%) of UK internet users claimed to use such services every week [9]. Previous studies have investigated user behaviour in a course-grained fashion in terms of item popularity, channel switching, and consumption rate. However, little effort has been spent to understand the contextual information related to user browsing activities; and in particular the behaviour that occurs before a user starts consuming an item of content. This initial ‘journey’ has a significant effect on user experience, resulting in them either finding something to watch or ultimately signing off from the service. Learning from these journeys could lead to significant advances in service personalisation and performance, such as dynamically optimising the user interface and strategically caching content based on user and time. Several online media services have become conscious to the fact that monitoring

and understanding user interactions are essential methods to improve efficiency of content distribution, advertisement revenues and quality of user experience. This is achieved through an analysis of mouse click sequences and content consumption characteristics.

In this study we provide an insight into such browsing habits. We analyse the browsing behaviour (including time of day) of 1,100+ users over the course of six months from a university campus online TV service that offers both live and VoD content to students and staff. Aiming at providing useful metrics to enhance recommenders and personalisation of media services, the paper applies clustering methods to generalise the findings in modelling user activities. We arrive at different ways in which users look for content and how behaviour patterns change over the 24 hours of the day.

The paper proceeds as follows. Section 2 provides background and discusses related work, Section 3 introduces the online TV system and the dataset analysed in this paper, Section 4 studies the temporal dynamics of user browsing behaviour through the clustering of user interactions, and Section 5 concludes and relates future work.

2. BACKGROUND AND RELATED WORK

The primary motivation of many studies (e.g. [1, 8]) has been to understand the impact of video requests on the network, driven by the impetus to improve performance through intelligent pre-fetching and content caching. Few studies have considered temporal dynamics: Yu et al. identified periodic changes in video requests (as a reflection of user interest) over different scales [14]. Cha et al. studied user behaviour in terms of channel switching and its relationship to content genre, time of day, and region [2]. Qiu et al. and Kim et al. incorporated diurnal patterns into modelling channel/genre popularity, and differentiated users into groups based on channel/genre interest [10, 6]. Sanchez et al. modelled user interests in an IPTV service based on content consumption [11]. Keller et al. also studied user interactions related to search to improve post-search navigation [5]. Wang et al. employed sequence- and time-based clickstreams to detect malicious users through clustering [13].

Current work on activity in multimedia systems is mostly based on one or two deterministic metrics such as playback requests and user ratings. While these metrics are effective in studying the one-mode user-content relationship for user-based recommenders, there is significant value in studying the contextual information capturing a user’s stream of browsing activities prior to the consumption of content. An example of this is the effect of how system features are pre-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

MM '14, November 03 – 07 2014, Orlando, FL, USA

Copyright 2014 ACM 978-1-4503-3063-3/14/11...\$15.00.

<http://dx.doi.org/10.1145/2647868.2654980>.

sented on content discovery and consumption (cf. [12]) as well as user experience.

The modelling of web page browsing activities is generally based on the frequency and order of clicks or page loads. For a media service, the amount of time spent on web pages of specific functions also provides crucial context indicating user’s content navigation preference such as relaxed browsing without specific intention or quick keyword search to locate a piece of content in mind.

This paper focuses on mapping user browsing journeys in an online TV service through defining and modelling the contextual information related to how the users find content.

3. VISION SERVICE AND USER DATA

3.1 IPTV system overview

Vision is an online TV service that is free for students and staff of Lancaster University. It provides a Living Lab environment with real users and workload for research and development work related to content retrieval, media distribution and quality of user experience. Between service launch in October 2013 and March 2014, Vision drew 1,100+ registered users with $\approx 34,710$ sessions. In accordance with the university’s ethical guidelines, Vision does *not* collect any personal user information such as age, gender, etc.

A total of 30 national TV and 20 radio channels are made available as live streams using a number of transcoding and media servers. The TV channels are continuously recorded and automatically partitioned into programmes, based on EPG data. These programmes are added to the VoD repository which holds over 30 days worth of catch-up content.

User facing functionalities are manifested in 7 web pages, each distinguishable by a unique ID: *Dashboard*, *What’s On*, *On Demand*, *Programme Guide*, *My Library*, *My History*, *Search Results*, and *Playback*. The Dashboard (Fig. 1(a)) incorporates a recommender (based on highly customised collaborative filtering), watch list (a shortened list of saved items), and social trending features (trending content within the user community). Playback page (Fig. 1(b)) encapsulates a media player, programme information, and social features. What’s On offers currently airing and upcoming programmes of each live TV channel. Users can browse the entire VoD repository via the On Demand page by name, genre, length, etc. Programme Guide provides a grid view of the TV schedule. Vision users may also add items to a personal library (via My Library). Users can access their viewing record on My History and resume playing a programme from where they left off on any device.

3.2 User and service statistics

As a research platform, Vision focuses on gathering user feedback through periodic survey and finely grained quantitative user activity data. A dedicated *statistics service* captures user interactions for comprehensive social and user behaviour analysis. A *report engine* resides in user devices (in JavaScript) to i) catch time-coded events such as playback requests, periodic heartbeats of media player, keyword searches or page loads, ii) pre-process the data, and iii) post the statistics to the *stats ingest service* to be stored for data analysis. An aggregation of faceted usage statistics reveals page navigation and content viewing experience in Vision. For instance, the system is able to tell that from 18:05:01, user 3 spent 2 minutes going through live programmes on




Figure 1: Vision web interface

What’s On, then spent 14 minutes browsing On Demand, and finally resumed an episode of *Flog It* from My History.

3.3 Journey of user interaction

We extract all recorded user activities of 1109 users (excluding admin and test accounts) for the period 02/10/13–17/03/14. A user session is identified when the user activity stream becomes idle for 30 minutes, the standard notion of a visit in Web analytics [7]. Each session may contain one or more playback requests. Every playback request is seen as the termination of a user’s journey of seeking content to watch. Therefore we define each playback request along with the relevant user interactions prior to the request as a *journey*. The journey gives an insight into how each user, at a particular time of the day, navigates between service pages before playback commences. A total of 67,089 journeys are recognised for the experimentation period. One specific element of user journey is the proportion of browsing time spent on each service page pt_n (where n is the page id). pt_n captures the characteristics of each user playback journey, hence is exploited as a metric to study the user similarities with respect to content navigation and to investigate whether browsing preference changes over time.

4. TEMPORAL DYNAMICS AND CLUSTERING OF BROWSING ACTIVITIES

4.1 Initial observations

Related work (such as [2, 1, 8]) has already shed light on how user playback requests change over different time scales (day, week, month, etc.) In this work, however, we investigate whether a similar time dependency applies to users’ browsing activities and their interactions with the service.

Our periodic user surveys and interviews suggested that most of our user population start using our TV service from the late afternoon after lectures. Many users also claimed that they have developed the habit of watching a few TV programmes before they go to sleep around midnight.

First, we examine feature use over the hours of the day. We find no strong deduction when using such co-variate time information. We conjecture that this is due to variance in user-days; e.g. two users could be using Vision while having breakfast, but that means 7am for one and 10am for the other. We then consider dividing the day into a number of time slots wherein user activity seems to conform to a strong trend. In Fig. 3, we plot the average number of all user journeys over the 24 hours. In order to achieve a




Figure 2: Violin plot of the average proportion of journey time spent on each page in each time segment.




Figure 3: Density of journeys per $1/4$ hour.

meaningful classification, we bin user activities into three main time segments: 5am-3pm, 3pm-11pm, and 11pm-5am.

Based on this, we inspect the journeys in each segment to investigate whether we can generalise the patterns of user behaviour in groups within each time segment and how the patterns change over time. Fig. 2 characterises our findings in a violin plot. For each of the seven Vision service pages per time segment, the figure offers a box plot to illustrate the median proportion of time spent on the page relative to the journey duration. The box plot also presents the first and third quartiles, signified by thin lines at the extremes. Additionally, the figure depicts the overall distribution as a density function plotted in colour around the box plot.

We learn some interesting findings from this figure, the first and most crucial of which is that the vast majority of users have journeys that belong to several clusters. Furthermore, there is a discernible difference in the role some of the pages play in user journeys across the time segments. For instance, My Library and Search are far more prevalent in the third time segment. This confirms the first of our suspicions that users utilise the system features in a different manner across various times in a day. However, the ‘hourglass’ like shapes of the densities suggest different features within as well as between time segments. Inspired by such findings, we launch a deeper analysis of the affinity user journeys have with the pages during the different time segments.

4.2 Analysis of browsing behaviour

We employ clustering to identify different types of users, expecting to see significant overlap between the different types of content browsing behaviour. Traditional nonparametric clustering methods such as K-means may struggle to identify any clusters [4]. Thus, we use model-based clustering in order to identify clusters and assume Gaussian distri-

butions for each cluster. Whilst the original page time pt_n data are proportions (restricted to $[0,1]$) we rescale these to the real line to make the Gaussian assumption more appropriate. All the calculations to fit the clusters were performed using the EMCluster [3] package in R.

The clustering of the journey data resulted in 4, 5, 5 clusters for the 5am-3pm, 3pm-11pm and 11pm-5am time segments, respectively, with the highest significance. The average proportion of time spent per page within each cluster is depicted in Fig. 4.

For the time period 5am-3pm we observe key differences between the 4 clusters. Out of the seven pages to seek content: cluster 1 browses in six pages with no usage of the My History; cluster 2 browses in two pages with 40% of the time spent on Dashboard and the remaining 60% spent browsing live channels on What’s On; cluster 3 browses in two pages with the majority of the time (86%) on Search; cluster 4 browses in all pages with 50% of time spent on My Library. In this time segment it is clear that we have two types of users who are easily satisfied, namely those in clusters 2 or 3 who either identify content through the Dashboard (which is the default landing page of the service) and through What’s On or Search, respectively. The remaining two groups, 1 and 4, are less easily satisfied and are willing to spend more time interacting with the entire system looking around all pages before deciding on content. These journeys also appear to be from users without a predetermined consumption intent, eventually finding content through one of the live channels or perhaps resorting to something from their profile (My Library/My History) when other means do not succeed.

For the 5 clusters of journeys in the time period 3pm-11pm we have the following key differences: Out of the seven pages to find content, cluster 1 uses two pages with most of the time (87%) spent on What’s On; cluster 2 uses three pages with 88% on Dashboard; cluster 3 uses all pages almost equally; cluster 4 uses four pages with 70% attention to On Demand; cluster 5 uses three pages with nearly all browsing activities (99%) on Search. Here, the majority of journeys (58.8%, cluster 1) favour live content. These are users who tune in at a specific time to watch primetime programmes. In contrast, users in cluster 5 identify content through keyword search. Interestingly, cluster 4 is the only cluster across the three time periods where browsing for On Demand content takes up the majority of their journey time.

For the 11pm-5am segment: cluster 1 uses four pages with 65% of time on What’s On; cluster 2 use five pages with 44%




Figure 4: The proportions of time spent on each page between clusters and across time segments.

on What’s On, 33% on My Library, 28% on On Demand; cluster 3 uses four pages with 47% Search and 36% Dashboard; cluster 4 uses all pages with 59% of time browsing Programme Guide; cluster 5 also uses all pages but with half of the time on My History. Unsurprisingly, browsing at this time of the day involves less time spent on What’s On as the quality of live content declines, and much more towards the different VoD channels. Most users find items from My Library or the On Demand catalogue. The next group is perhaps more specific and rely heavily on Search. The third group is more reliant on the Programme Guide as a means of catching up on programmes missed earlier in the day.

These results show clear differences between the time segments especially with access to VoD content (via On Demand, My Library, etc.) being common to several clusters in the 11pm-5am segment compared to belonging to 1 cluster in the other time segments. Similarly, the 11pm-5am segment is the only time when the Programme Guide is used as the primary content portal for a cluster. In contrast, the use of My Library is only prevalent in the 5am-3pm segment.

In summary, the results provide insight into the different habits of users browsing an online TV service. In particular, through the use of clustering we have identified distinctive groups of users with consistent browsing behaviour exhibited during different periods of the day.

5. CONCLUSIONS

We used model-based clustering to distinguish groups of similar user browsing habits in an online TV Living Lab environment. We identified a number of statistically significant clusters each of which capture thousands of user browsing journeys with distinct behavioural patterns. This could easily be leveraged by service features to personalise user experience in finding content (e.g. adapting the user interface, seeding personalised recommenders) or to increase revenue (e.g. promoting content).

Although our dataset relates to one online TV service, the followed methodology and derived conclusions can be adopted by similar online media services. The outcomes from studying this case study presents an initial step towards understanding user behaviour and content discovery in relation to different contextual factors. We plan future efforts to enhance the analysis with further contextual information. A metric we plan to include is the proportion of content watched beyond the end of a journey, which we

hope will provide insight into the relationship between user engagement and journey characteristics.

6. ACKNOWLEDGMENTS

This work was supported in part by European Commission FP7 grants 318343 (STEER) and 603662 (FI-Content2).

7. REFERENCES

- [1] H. Abrahamsson and M. Nordmark. Program popularity and viewer behaviour in a large TV-on-demand system. In *Proc. IMC*, pages 199–210. ACM, 2012.
- [2] M. Cha, P. Rodriguez, J. Crowcroft, S. Moon, and X. Amatriain. Watching television over an IP network. In *Proc. IMC*, pages 71–84. ACM, 2008.
- [3] W.-C. Chen, R. Maitra, and V. Melnykov. *EMCluster: EM Algorithm for Model-Based Clustering of Finite Mixture Gaussian Distribution*. R package v0.2-4.
- [4] K. A. Heller and Z. Ghahramani. A nonparametric bayesian approach to modeling overlapping clusters. In *Proc. Conference on Artificial Intelligence and Statistics*, pages 187–194, 2007.
- [5] M. Keller, P. Mühlshlegel, and H. Hartenstein. Search result presentation: Supporting post-search navigation by integration of taxonomy data. In *Proc. WWW Companion Conference*, pages 1269–1274, 2013.
- [6] J. Kim, C. Hwang, E. Paik, and Y. Lee. Analysis of IPTV user behaviors with MapReduce. In *Proc. Conference on Advanced Communication Technology*, pages 1199–1204, 2012.
- [7] R. Krishnan, S.S. and Sitaraman. Video stream quality impacts viewer behavior: Inferring causality using quasi-experimental designs. *IEEE/ACM Trans. Networking*, 21:2001–2014, 2013.
- [8] G. Nencioni, N. Sastry, J. Chandaria, and J. Crowcroft. Understanding and decreasing the network footprint of over-the-top on-demand delivery of TV content. In *Proc. WWW*, pages 965–976, 2013.
- [9] Ofcom. The communications market. <http://stakeholders.ofcom.org.uk/market-data-research/market-data/communications-market-reports/cmr13>, Aug 2013.
- [10] T. Qiu, Z. Ge, S. Lee, J. Wang, Q. Zhao, and J. Xu. Modeling channel popularity dynamics in a large IPTV system. In *Proc. SIGMETRICS*, pages 275–286. ACM, 2009.
- [11] F. Sanchez, M. Barrilero, F. Alvarez, and G. Cisneros. User interest modeling for social TV-recommender systems based on audiovisual consumption. *Multimedia Systems*, 19:493–507, 2013.
- [12] G. Tyson, Y. Elkhatib, N. Sastry, and S. Uhlig. Demystifying porn 2.0: A look into a major adult video streaming website. In *Proc. IMC*, pages 417–426. ACM, 2013.
- [13] G. Wang, T. Konolige, C. Wilson, X. Wang, H. Zheng, and B. Y. Zhao. You are how you click: Clickstream analysis for sybil detection. In *Proc. USENIX Security Symposium*, pages 241–256, 2013.
- [14] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng. Understanding user behavior in large-scale video-on-demand systems. *SIGOPS Operating System Review*, 40(4):333–344, 2006.