

YouTube Monetization: Creating User-Centric Experiences Using Large Scale Data

Dr. Ramesh R. Sarukkai
YouTube/Google Inc
1600 Amphitheatre Parkway
Mountain View
CA-94043, USA
sarukkai@google.com

ABSTRACT

Over the last 4 years, YouTube has grown from a viral video sharing site to a platform that fuels a win-win ecosystem for video content creators, advertisers and users. A key driving force behind this successful transformation is building out products/platforms that focus and optimize for the user. In this talk, we focus on some illustrative user-centric efforts such as user-controlled skippable ads (TrueView Instream) and dynamic ad loads (data driven system that balances user impact with advertising opportunities). Leveraging very large amounts of real-time activity data is paramount to successfully building and deploying such user-centric models. We conclude the talk with challenges and opportunities in this important area of real-time user analysis and large data modeling.

Categories and Subject Descriptors

H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval—Information filtering; I.2.6 [Artificial Intelligence]: Learning—Parameter learning.

Keywords

Real-time analysis, user modeling, advertising, optimization, large scale data.

1. INTRODUCTION

The internet has seen tremendous opportunity and growth over the past two decades – transforming from an information search system to a vital cloud infrastructure that powers a variety of different information needs, social activities and interactive services. With such continued usage of web powered services and activities, we have the opportunity to optimize experiences that for our users, even more effectively than before. Specifically, we are moving from an era of comprehensiveness to an era of precision for our users – find the right information, content or action for our users, at the right time, in the right context.

For YouTube, this journey has been a transformational one – going from a viral video site to a platform that creates a win-win ecosystem for our content uploaders, advertisers and users. One of the benefits of advertising on YouTube is that we provide users free access to interesting video content, while creating a viable ecosystem whereby content creators benefit from the advertising revenue. In this talk, we discuss how we create user centric experiences around advertising – namely TrueView Instream and Dynamic Ad Loads.

2. User-Centric Experiences

The prime motivation that we centered monetization experiences around is the user. Specifically, we ran experiments that allow us to precisely measure user impact due to ads – an example would be looking at user impact measured by video watch time on site as a proxy. There are different aspects to such user impact – user state (user has insufficient time, or is focused on their activity at that point), ads quality (how relevant is the ad), and ads latency (does the experience slow down due to ads).

Putting users first meant looking at ways in which they can engage with ads when they want but otherwise continue uninterrupted with their video watching. The first concept that helped steer us towards that goal was the notion of skippable ads – ads that are built on the notion of CPV (Cost Per View). For users, the benefit is more control over their experience, and for advertisers the benefit is that they pay for engaged views. Using large scale data modeling, our advertising systems optimize the ad selection to show ads effectively to users.

Another problem that we applied large scale data and user modeling is dynamic ad loads: deciding when to show an in-stream advertisement? The traditional approach of pre-defined breaks is not effective on the web, especially for short-form content, both due to the dynamic nature of video consumption as well as the diversity of content. An alternate approach to this problem is to take into account users active state and real-time interactions with the site, holistically rather than treat each video view and advertising event in isolation. This dynamic advertising load system (DAL) uses per visit user data and state to train models that predict user response to ads. The tradeoff here involves balancing user happiness (as measured by predicted time spent watching a video) versus the revenue opportunity for our content creators, in a unified manner modeled with a joint cost function. The net outcome was a win-win situation, whereby we are able to optimize both for the needs of the user experience (as measured by time spent watching videos) and the overall advertising revenue.

3. CONCLUSION

Optimizing our products and services towards user centric metrics is paramount to the future success of internet and media – we are in an era where user time is the most valuable commodity. Leveraging big data and modeling to tailor web applications and ads to balance such user satisfaction metrics with product/revenue goals allows us to improve the overall satisfaction for users – a notion we call “user-centric experiences”. We presented two illustrative efforts: the first is transforming instream advertising experiences to a skippable performance model; the second effort is the Dynamic Ad Loads system to optimize for user watch time with advertising opportunity. We believe this direction is fertile with latent ideas and potential applications to many other domains in the web.