

Introduction

Buzzfeed’s current strategy is focused on user growth – acquisition and retention of users. Customer acquisition ensures growth, but that growth is only meaningful if we can continue engaging the new and existing user base. To do so organically, we need to approach growth from a mindset of wanting to help users reach their goals – whether it is to be entertained or educated.

In this proposal, I focus on push notifications as a growth instrument for BuzzFeed. I propose a vision of a well-constructed push notification system, as well as instrumentation around it to gain quality user insight and measures of user experience.

Framework of values

Push notification acts as a door to the app experience, a call to its adventure. If it is not opened, there is no experience, which is why it needs to be the right door at the right time. You cannot force users to open it – it is a decision they have ultimate control over. Moreover, an average user gets 46 notifications per day [1]; anytime an app does not make valuable use of user’s attention, it gets one step closer to its permanent silence – “mute notifications”.

To align ourselves with user expectations, we need to understand them first. A great push notification needs to be **timely and relevant**. It should help experienced users effectively navigate the content they want to enjoy, while introducing new users gently into what is available to them. A great push notification system is an assistant that knows when you want to read what content.

Value	Description
Timely	We want the app to sync with user’s routines (e.g. commute schedule), so we don’t interrupt them. Also, we want to push appropriate content at appropriate time of day, day of the week, date.
Relevant	We want to use push notifications to help users consume content that is valuable to them.

Table 1: Push notification values user values.

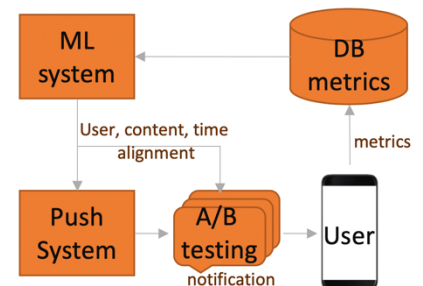
Under this framework of values, I propose metrics to measure user experience and its change as I try to improve it.

Current Experience

After using the app for a week, my experience as a new user has been pleasantly non-invasive, but too vague to associate strongly with the app and how I would navigate it. I was recommended content that seemed to follow general trends at the time (3 iPhone related, 2 movie related), but did not cater to my interests I tried to show through the app by searching and reading about camping and Halloween costumes.

Methodology

I propose our first-party push notification instrumentation in 3 parts: metrics of performance and how they align with our framework of values in Table 1, a machine learning framework that uses those metrics to help align users with their desired content at the right time, an A/B testing methodology that allows us to track defined metrics as we optimize user experience.



Methodology: Metrics of performance and their usage.

I focus on metrics related directly to push notifications. User engagement with content is outside the scope of push notifications, so I refer to those metrics as **session_engagement** (e.g. average time spent on page, articles opened, scroll-depth). It is important to mention that session_engagement metrics should include signals of user fatigue so that our optimization does not pursue to exploit users’ bounded rationality, but rather their genuine interest.

I distinguish between two event types: “Swipe” (close notification) and “Tap” (open notification).

Metric	Description
timestamp_of_action	Time the action was taken.
timestamp_of_notification	Time push notification got posted.
action_delay	Delay between timestamp_of_notification and timestamp_of_action.
timezone	Timezone of the device.
time_of_day	User perceived time of day.
user_id	Uniquely identifying user id.
content_id	Uniquely identifying content id.
device_type	What type of device (e.g. mobile, tablet, OS) is the user on.
notification_opened	Was notification open or closed.
session_from_notification	(if "Tap") Propagate this flag to engagement metrics to distinguish session information collected from regular use vs. notification driven.


Table 2: Tap/Swipe notification event schema.

- Session_engagement vs. time_of_day, action_delay tell us about **timeliness** of notifications.
- Knowing certain sessions came from a push notification rather than regular app browsing, we can compare if our recommended content is more engaging across session_engagement metrics than a regular session would be. This tells us about **relevancy** of notifications.
- If we see a significant change in *open ratio* (weekly average of notification_opened), it implies a change in **user perception** of our notifications. If it increases, we are on the right track; if it drops, we should be cautious of doing something harmful; if it is suddenly zero, it means the user likely muted the app and we should save their notification history as a negative example.

Methodology: Machine learning instrumentation

Buzzfeed publishes about 222 articles per day and has over 200M unique monthly visitors [2,3]. To effectively navigate such scale, we have to rely on automatic classification of content and users. Metrics that we collected allow us to extract such information through machine learning (e.g. collaborative filtering) techniques. Machine learning could give us a quantitative way to measure if particular users prefer specific time of the day for different content. We have 3 main dimensions across which we can find patterns:

- Users vs Content: Does this user prefer quizzes to share with coworkers or extensive articles to stay up to date with the world?
- Users vs. Time: Is this user an "evening commuter", "morning coffee briefing", "after-work unwind" type of person?
- Content vs. Time: Is this type of content best for particular time of the day?



This allows us to tune push notifications to user's preferences automatically and at scale. However, an unwanted side-effect of many personalization techniques is their convergence into predictably similar and familiar pockets of experience. Novelty is just as important, and it is a tradeoff that is tunable through experimentation and explicit metrics of user fulfillment that fall under session_engagement.

Methodology: A/B testing

I propose an A/B testing tool to track proposed metrics at individual user level, to maximize the quality of their experience. The A/B testing framework allows us to attribute changes in performance to tangible decisions made under controlled experiment conditions. In particular, we can optimize the timing and frequency of push notifications for certain content and user types that we identified via ML instrumentation. Such tool could help suggest significant user group size for a chosen experiment, length of the experiment to see novelty effects fade away, user groups that are relevant and do not overlap in their ongoing experiments.

[1] [businessofapps](#) "Push Notifications Statistics (2019)"

[2] BuzzFeed

[3] The Atlantic. "How many stories do newspapers publish per day."