

# Fun with ML, Stream Analytics and PowerBI - Observing Virality in Real Time

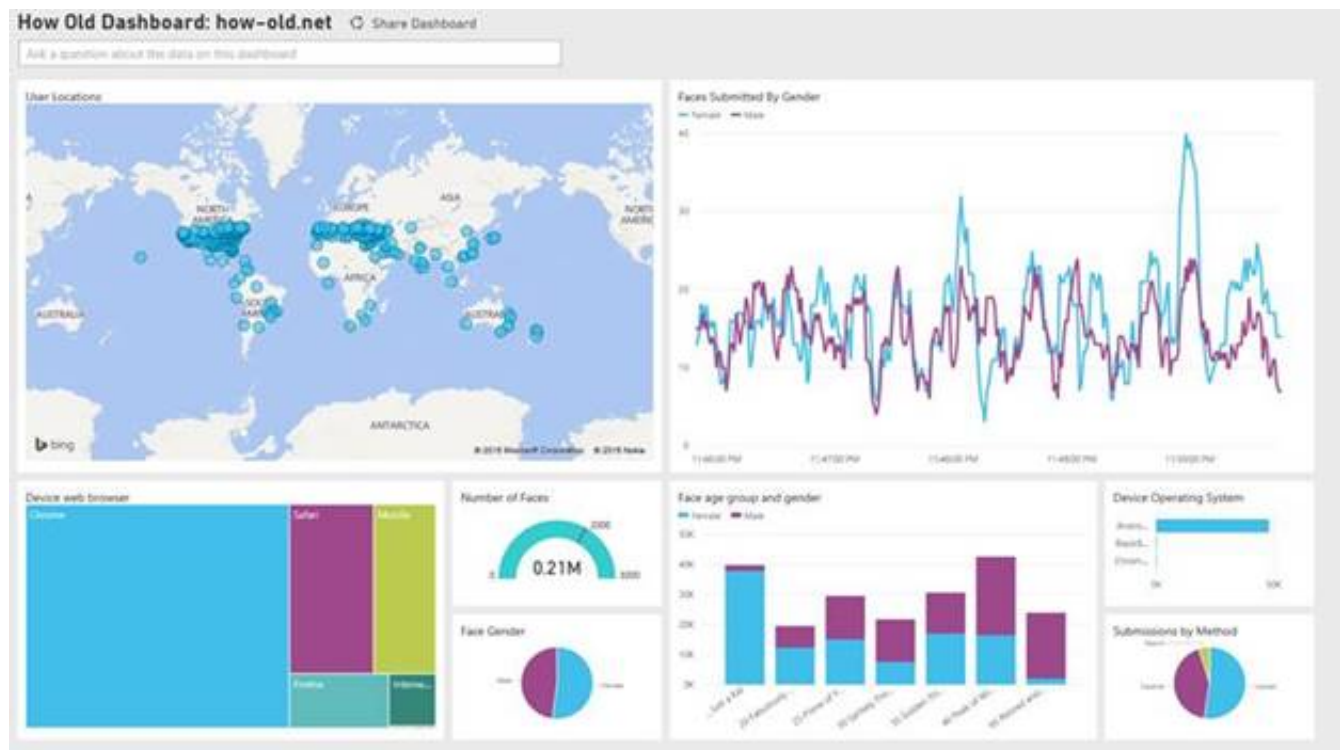
*This post is authored by Corom Thompson and Santosh Balasubramanian, Engineers in Information Management and Machine Learning at Microsoft*

Updated 5/2/2015

*We've had some questions so we updated this post to be more clear. To answer the top one: No we don't store photos, we don't share them and we only use them to guess your age and gender. The photos are discarded from memory once we guess. While we use the terms of service very common in our industry, and similar to most other online services, we have chosen not to store or use the photos in any way other than to temporarily process them to guess your age.*

This is a fun story of how we were expecting perhaps 50 users for a test but - in the end - got over 35,000 users and saw the whole thing unfold in real time.

We were building a demo for [the day 2 keynote](#) of Microsoft's [Build2015 developer conference](#). We wanted to showcase how developers can easily and quickly build intelligent applications using Azure services. Using our newly released [Face detection API's](#) we set up an age guessing website called <http://how-old.net> on Azure. This page lets users upload a picture and have the API predict the age and gender of any faces recognized in that picture. Now, while the API is reasonably good at locating the faces and identifying gender it isn't particularly accurate with age, but it's often good for a laugh and users have fun with it. We sent email to a group of several hundred within Microsoft asking them to try the page for a few minutes and give us feedback - optimistically hoping for a few tens of people to try it out and generate some usage data to test the demo. But here is what our dashboard showed in three hours:



Within hours, over 210,000 images had been submitted and we had 35000 users from all over the world (about 29k of them from Turkey, as it turned out - apparently there were a bunch of tweets from Turkey mentioning this page).

Unique count of users by country or region

COUNTRY OR REGION	UNIQUE COUNT
Turkey	29 K
United States	5.6 K
United Kingdom	885
Germany	555
Canada	442

The demo showed real time insights about how people were using this tool. For instance, we had assumed that folks would mostly select from pre-canned images or use the Bing image search box on the page. But over half the pictures analyzed were of people uploading their own images. This insight prompted us to improve the user experience and we did some additional testing around image uploads from mobile devices.

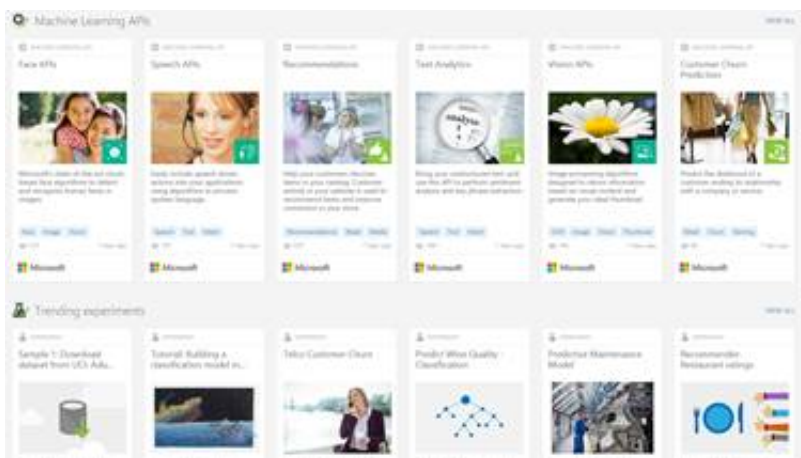
## So What's The Magic Behind All This?

This may be hard to believe but it took a couple of developers just a day to put this whole solution together, from the web page to the Machine Learning APIs to the real time streaming analytics and real time BI. This turned out to be a great example of the agility and creative power Azure developers get. The key components of this application are:

- Extracting the gender and age of the people in these pictures.
- Obtaining real time insights on the data extracted above.
- Creating real time dashboards to view the above results.

## Extracting Gender and Age

We wanted to create an experience that was intelligent and fun could capture the attention of people globally, so we looked at the APIs available in the [Azure Machine Learning Gallery](#). The gallery contains many finished intelligent services for [Face](#), [Speech](#), and [Vision](#) which are part of a new suite called [Project Oxford](#) from Bing and Microsoft Research. The [Face API has a demo page](#) that uses the API to detect and extract information about faces in a photograph. We found the ability of the face API to estimate age and gender to be particularly interesting and chose this aspect of it for our project. To make the experience more fun we used the face API alongside the [Bing Search API](#) from the Azure marketplace to create <http://how-old.net>.



In addition to age and gender, we also used additional information provided by standard web browsers, such as the User Agent string that comes with every standard HTTP call and the latitude and longitude of location from where the picture was uploaded. These can be used to calculate standard website usage statistics such as the number of hits from iPhones, Windows or Android, or places where how-old.net is most popular. This is represented in following JSON document:

```
[ { "event_datetime": "2015-04-27T01:48:41.5852923Z",
```

```

"user_id": "91539922310b4f468c3f76de08b15416", "session_id": "fbb8b522-6a2b-457b-bc86-62e286045452",
"submission_method": "Search",
"face": { "age": 23.0, "gender": "Female" },
"location_city": { "latitude": 47.6, "longitude": -122.3 },
"is_mobile_device": true, "browser_type": "Safari", "platform": "iOS", "mobile_device_model": "iPhone"
}

```

## Real Time Insights

To understand the patterns in the real time data from this website, we used a set of new Microsoft Azure streaming services.

We brought in the data using [Azure Event Hubs](#), a highly scalable publish-subscribe ingestor that can intake millions of events per second. We use the Event Hubs API to stream the JSON document from the web page when the user uploads a picture. Note that **the photo is not saved** and no information identifying or linking to users is saved (we have no emails or logins or usernames), only the JSON document is streamed to Azure Event Hubs.

Next we needed a stream processing service to aggregate and process the information from thousands of users uploading pictures in real time. For this we use [Azure Stream Analytics](#) (ASA), a fully managed low latency high throughput stream processing solution. ASA lets you write your stream processing logic in a very simple SQL - like language.




An example of using ASA, if you want to get the count of "gender" in a 10 second window with a result written every second, all you need is a very simple query to aggregate this information:

```

SELECT
  System.Timestamp AS OutTime,
  Face.gender AS Gender,
  Count(*) AS Count
FROM
  StreamInput
GROUP BY HoppingWindow(second,10,1),
  Face.gender

```



In the above query, we are selecting the Time when the result is written (OutTime), Gender, and count of gender. StreamInput is the alias of the Event Hub to which the streaming log data is sent. This is done in a hopping window of 10 seconds, with a hop of 1 second. This query gives the aggregate count of Female and Male faces in the uploaded pictures and this information can be displayed in a dashboard. You can have multiple stream processing queries on data coming from the same Event Hub.

## Real Time Dashboards

We use PowerBI to display the results in a real time dashboard. All we did was to choose PowerBI as the output of our stream analytics job (click here to [learn how](#)). Then we went to <http://www.powerbi.com>, and selected the dataset and table created by ASA. There is no additional coding needed to create real time dashboards.



In this example we have a couple of stream analytics queries. One aggregates age into an age range and passes in other fields such as location to PowerBI, and the other is the query mentioned above. PowerBI lets you easily create a variety of visualizations including maps, line charts, tree view charts and more. Charts get updated in real time as data is generated by users uploading pictures at <http://how-old.net>. Additionally you can ask natural language questions too (e.g. "What's the count of people using IOS by gender by age group") and the charts that are displayed as a result of such a question can be pinned to the real time dashboard.

Go try out <http://how-old.net> for yourself ([#HowOldRobot](#)) - we hope you have fun with it and are inspired to

create your own applications using Azure services and the APIs available in the [ML Gallery](#).

Corom and Santosh