

The International Affiliation Network of YouTube Trends

Edward L. Platt and Rahul Bhargava and Ethan Zuckerman

MIT Center for Civic Media

{elplatt, rahulb, ethanz}@media.mit.edu

Abstract

Online video, a ubiquitous, visual, and highly shareable medium, is well-suited to crossing geographic, cultural, and linguistic barriers. Trending videos in particular, by virtue of reaching a large number of viewers in a short span of time, are powerful as both influencers and indicators of international communication flows. In this work, we study a large set of videos trending across 57 nations, collected from YouTube over a 7-month period. We consider the set as a network of content flowing between nations, then develop conditional co-affiliation, a nation-nation co-affiliation index that enables a meaningful interpretation of network path length and the application of betweenness centrality. We observe a highly-interlinked network with remarkably similar co-affiliation levels between very different nations. However, Arabic-speaking nations appear more isolated, with the U.A.E. emerging as a key bridge. By analyzing video trend lifespans, we show that nations having many globally-popular video trends are reliably not the nation where those trends are strongest: we see no evidence to support the widely discussed idea of cultural exporter or trendsetter nations. We model correlations between co-affiliation and a selection of contextual factors. We note a surprisingly complex interaction between migration and shared video trends. Consistent with existing work on video popularity, we find that long trending times within one nation do not necessarily translate to reaching a wide global audience. This work expands on previous studies of the geographic popularity of videos by incorporating trending data and extending our analysis from video-nation affiliations to nation-nation co-affiliations. Characterizing these relationships is key to understanding the international cultural impact and potential of online video.

Introduction

As access to digital communication becomes increasingly common around the world, it becomes easier to exchange ideas between cultures, crossing once-prohibitive geographic distances and national boundaries. Online video is particularly intriguing as a potential vector for cultural communication. As a visual medium, video offers the potential to cross linguistic and literacy barriers. Video is also able to capture cultural experiences such as dance and music, which are difficult to convey through text-based media.

Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

But are new communication technologies actually being used to share ideas globally, or are they simply reflecting preexisting cultural channels? And how do cultural, political, and geographic factors influence international communication? By analyzing usage data from digital communications platforms, we can begin to answer these questions. In this paper, we focus on trending data from the YouTube video sharing platform to examine the international usage of online video.

The study of changing communication technology and its influence on international culture has a long history, preceding the advent of digital communication. Norris and Inglehart review this history and present four distinct hypotheses (Norris and Inglehart 2009). The *L.A. effect* (referring to Hollywood’s export of culture through film) proposes that increased cross-cultural communication has a homogenizing effect, and threatens to erase distinct cultural identities by replacing them with a Western culture. In contrast, the *Bangalore effect* suggests that communication enables a two-way cultural fusion in which cultures borrow, remix, and share ideas while maintaining distinct identities. A third hypothesis, the *Taliban effect*, describes a backlash against outside cultural influences, with increased communication polarizing cultures. Norris and Inglehart also put forth their own proposal, the *Firewall effect*, suggesting that even with geographic barriers removed, societal and individual barriers such as trade restrictions and differing cultural values still pose a significant obstacle to the flow of culture. (Zuckerman 2013) has applied these models to the flow of ideas online, arguing that although technology can cross cultural boundaries, it does not do so automatically. In particular, the communication of ideas across cultures requires both *xenophiles* who enjoy exploring other cultures and *bridge figures* who, with deep understandings of multiple cultures, can facilitate that exploration.

In this paper, we use geolocated online video trending data from YouTube Trends to construct and investigate a network of nations linked by shared video trends. YouTube defines trending videos as those that “have become popular because they were embedded in the web’s most popular websites and a significant number of people viewed the video externally in addition to on youtube.com.” (YouTube 2015) YouTube further divides video trends into “most shared” and “most viewed.” For this project we consider the “most

viewed” trends. Trending videos are displayed prominently on YouTube’s homepage, further increasing their popularity. For this project, we consider nations, rather than individuals, to be our fundamental unit of analysis. While nations may have different populations and total video views, YouTube displays the same number of trending videos to users in all nations, and we concern ourselves only with the fraction of those videos each pair of nations shares.

We focus primarily on whether the trending videos displayed to users in different nations cross existing cultural and political boundaries. We caution against interpreting our results as reflecting intrinsic properties of national cultures or international relationships. Trending videos do provide a view into user behavior, but YouTube users may not be representative of entire national populations. A nation’s trending videos are influenced by YouTube’s algorithm, which is undisclosed and may differ over time or place. Even with a known, consistent algorithm, the notion of *algorithmic objectivity* is a fiction (Gillespie 2014). Similarly, we limit the scope of this project to trending videos only, rather than all YouTube videos. Regardless, trending videos have the potential to expose different nations to common ideas. Our purpose here is to identify whether, and under what circumstances, that potential is realized in the context of trending videos.

We begin by using trending video data to construct an international network of video trend co-affiliations. We analyze the network properties and identify key hubs and bridges. We apply hierarchical clustering to find communities of strongly-connected nations. We also model the influence of cultural, political, and geographic factors on shared video trends, focusing on *contextual factors* describing the nation rather than *event-oriented* factors (Chang and Lee 1992).

Our major contributions and findings are:

Conditional co-affiliation: We develop a co-affiliation index based on conditional probability. This index allows meaningful comparisons of multi-edge path lengths, enabling the application of metrics such as betweenness centrality.

Trends are cosmopolitan: Most nations share about the same number of trending videos with each other nation, regardless of cultural, geographic, and political differences. We see a notable exception in Arabic-speaking nations, which share significantly more trends with each other than with other nations. The U.A.E. emerges as a key bridge between these nations and the rest.

No global trendsetters: When a nation’s trends are shared with a larger number of other nations, those trends tend to be strongest in another nation. We see no evidence of trendsetter nations with many far-reaching, locally-based trends.

High migration means less global: Surprisingly, nations with higher migrant stock tend to share fewer trends with other nations in general. However, bilateral migration between a pair of nations does correlate with more trending videos in common.

Different models for different factors: We find that, of the four models of communication we consider, no single model is consistent with all contextual factors, but each model is consistent with some factors. This finding suggests a multifaceted relationship between national contextual factors and video trends.

Related Work

Considerable research has already examined the viewing habits of YouTube users. The popularity of YouTube videos have been shown to follow a non-Zipf, power-law distribution with a cutoff (Gill et al. 2007; Cheng, Dale, and Liu 2008; Cha et al. 2009) with videos’ active lives following a Pareto distribution (Cheng, Dale, and Liu 2008). (Brodersen, Scellato, and Wattenhofer 2012) found that YouTube videos are highly local, with 50% of videos having 70% of their views from a single nation. The spread of videos has been shown to occur in stages, with social sharing, subscription, and search driving popularity at early stages while non-social, centralized channels drive popularity in later stages (Brodersen, Scellato, and Wattenhofer 2012; Susarla, Oh, and Tan 2012).

Recent research has also looked at how digital technology can be used to map and understand international culture. Using a large international email corpus, (State et al. 2013) investigated correlations between international affinities and economic, political, and cultural factors. They found that the international affinities reflected by email traffic were consistent with the civilizations proposed by Huntington, and that affinities were correlated to economic disparity, consistent with World Systems Theory.

By analyzing bilingual Wikipedia and Twitter users as well as book translation data, (Ronen 2013) constructed a global language network that reveals the possible flow of novel information through translation. Within this network, English was identified as a significant hub language, used to pass information between other, less central languages. Spanish, German, French, Russian, Malay, and Portuguese acted similarly, but were somewhat less central.

(Taneja and Wu 2014) used website co-visitation data to argue that most web traffic is local, and that China’s isolation from the global Internet is less a function of government censorship and more a result of linguistic and cultural isolation. They see similar patterns of national and linguistic isolation in other nations where government censorship of the Internet is not a significant factor. Their work suggests that the transmission of ideas across cultures is challenging due to the lack of common cultural spaces, and we believe YouTube may be one of these spaces.

Methods and Data

Our analysis is based on data collected from YouTube Trends between May 15, 2013 and December 23, 2013. Once per day, we sampled the identifiers of the top 10 trending videos in the “most viewed” category for each available nation. YouTube does not disclose the update frequency of the list, but we observe it to be typically less than once per hour. To reduce bias from incomplete data, we consider a

subset of nations and days having complete data. Our final data set includes 36,274 unique trending videos, over a 169 day period, across 57 nations.

For the purposes of this paper, we model each nation as a multiset, or “bag of videos.” Videos are counted once for each day they appear in a nation’s top 10 trending list. In the interest of simplicity, we do not consider the timing and ordering of videos. This data can be represented by a video-nation matrix B , where $B_{v,c}$ is the number of times video v trends in nation c , analogous to the document-term matrix commonly used in natural language processing.

The video-nation matrix B can be interpreted as a weighted affiliation network (Borgatti and Halgin 2011). The weight of a video-nation affiliation is the number of times the video trends in that nation. We can also represent the video-network affiliation network as a bipartite graph with B as the biadjacency matrix.

Quantifying Nation-Nation Similarity

To compare nations directly, we construct a weighted, unipartite, nation-nation co-affiliation matrix A from the bipartite affiliation matrix B . In this case, “co-affiliation” simply means a similarity index derived from common video affiliations. Many methods exist to calculate co-affiliation, including the well-known Pearson correlation, multiplication of the biadjacency matrix by its transpose (Breiger 1974), the Jaccard/Tanimoto coefficient (Tanimoto 1957), the Bhattacharyya coefficient (Bhattacharyya 1943), and the Bonacich co-affiliation (Bonacich 1972). We opt for a probabilistic approach and use a generalized form of the “exposure” coefficient of (Ronen 2013), which we call *conditional co-affiliation*. This quantity has several novel properties, which we develop here.

We define the conditional co-affiliation $e_{j,k}$ as the conditional probability that a video trends in nation j , given that it trends in nation k :

$$e_{j,k} = \Pr(J|K) = \frac{\sum_{v \in J \cap K} B_{v,k}}{\sum_{v \in K} B_{v,k}}, \quad (1)$$

where J and K represent the multisets of all videos which trend in nations j and k .

We note that the $e_{j,k} \neq e_{k,j}$ and that the resulting co-affiliation network is, in general, directed. We can define a symmetrized conditional co-affiliation $s_{j,k}$ equal to the probability that a video trends in both countries, given that it trends in either:

$$s_{j,k} = \Pr(J \cap K | J \cup K) \quad (2)$$

$$= \frac{\sum_{v \in J \cap K} (B_{v,j} + B_{v,k})}{\sum_{v \in J} B_{v,j} + \sum_{v \in K} B_{v,k}}. \quad (3)$$

We proceed using the symmetric version, and will refer to it simply as co-affiliation throughout this paper.

An analogous measure of dissimilarity can be defined as the self-information of the conditional co-affiliation:

$$I_{j,k} = -\log_2 s_{j,k}. \quad (4)$$

The self-information decreases monotonically from infinity to 0 as the conditional co-affiliation increases from 0 to 1.

We use both depending on whether a measure of similarity or dissimilarity is needed.

In addition to providing an interpretation for the edge weights of the co-affiliation matrix, the conditional co-affiliation provides a meaningful interpretations of paths. Consider the probability of a trend occurring in nations J and K given that it occurs in L :

$$\Pr(J \cap K | L) = \Pr(J|K \cap L) \Pr(K|L). \quad (5)$$

With the simplifying assumption that the probabilities of a trend occurring in J and K are conditionally independent, (5) becomes:

$$\begin{aligned} \Pr(J \cap K | L) &= \Pr(J|K) \Pr(K|L) \quad (6) \\ &= e_{j,k} e_{k,l}. \quad (7) \end{aligned}$$

In other words, multiplying edge weights along a path gives the probability that a video trends in every nation along the path, given that it trends in the nation at the path’s beginning. The value $\Pr(J \cap K | L)$ can be compared to the value of $\Pr(J|L)$ to determine if nations j and l are more likely to share trends directly, or through an intermediary nation k . Similarly, the symmetric $s_{j,k}$ can be seen as the probability a trend will span the divide between two nations. The product $s_{j,k} s_{k,l}$ represents the probability of a trend spanning both pairs of nations, once again assuming independence.

If we instead consider edge weights given by the self-information, the products becomes sums, a more conventional way to combine edge weights into a path length. For this reason, we use self-information for calculations that involve path lengths, such as betweenness centrality.

For consistency, we also apply the conditional co-affiliation to the contextual data used in this paper. For example, we measure the strength of migration between two nations by calculating the probability that a person chosen at random from two nations will have been born in one and live in the other.

Bilateral and unilateral factors

We correlate video trend co-affiliation to contextual factors by applying ordinary least squares linear regression. Our observations are the co-affiliation values of nation pairs, and thus bilateral in nature. However we include both bilateral and unilateral factors as independent variables in the regression. For each bilateral factor, such as whether two nations share a major language, our hypotheses are simply that co-affiliation is positively or negatively correlated.

For unilateral factors such as GDP, there are two values per observation, one for each nation in the pair. We require a way to include these unilateral factors in a model of the intrinsically bilateral co-affiliation. To do so we include two independent variables for each unilateral factor: one for the nation with the lower value, and one for the higher. Each of these two values can be positively or negatively correlated with co-affiliation, giving four possible hypotheses.

Properties of Trending Videos

We characterize each video’s trending behavior according to how many nations that trend reaches and how many days it lasts. Specifically, we define the following parameters:

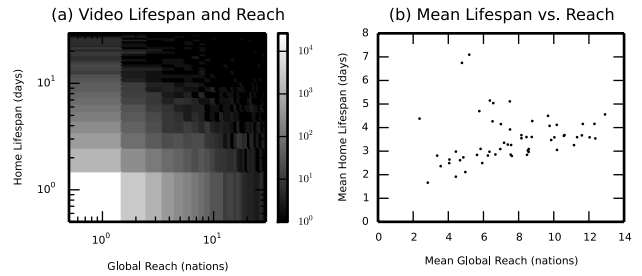


Figure 2: Video home lifespan and global reach.

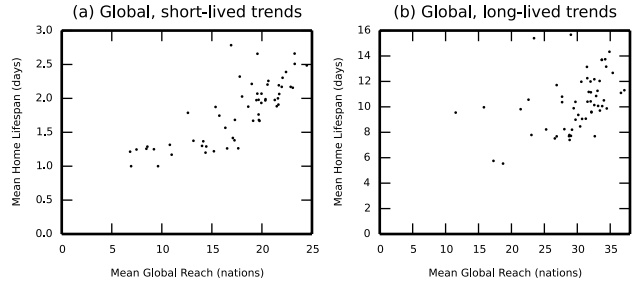


Figure 3: These plots show national means when considering only (a) short-lived and (b) long-lived video trends, relative to overall national means. National mean lifespan and reach show correlations of $r = 0.79$ ($p < 0.001$) and $r = 0.41$ ($p = 0.002$) when considering short-lived and long-lived trends, respectively.

Figure 1: Fitted lines correspond to the maximum-likelihood Zeta (discrete Pareto) distribution. Lifespan and global reach exhibit shape parameters of $\lambda_L = -3.13$ and $\lambda_R = -2.71$, respectively.

Global reach: The number of nations a video trends in.

Local lifespan: The number of days between a videos first and last appearance (inclusive) in a specific nation.

Home lifespan: The maximum local lifespan of a video in any nation.

Lifespan ratio: The ratio of the local lifespan to the home lifespan.

We refer to the longest single-nation lifespan as the home lifespan based on the observation that video popularity tends to start in a single geographic region, expand to new regions, and then contract back to the original region (Brodersen, Scellato, and Wattenhofer 2012). By comparing the lifespan of a video in a nation to the home lifespan, we get a rough estimate of how close a nation is to a trend’s origin (not necessarily the same as the video’s origin). When a nation’s lifespan ratio is 1, we consider it to be the geographic home of a video trend (not necessarily the same as the nation that produced the video). The lifespan ratio will decrease for farther-removed nations. A nation’s average lifespan ratio represents the introversion of a nation: how often that nation participates in internal vs external trends.

We find that, on average, videos trend for a lifespan of 1.88 days and reach 1.62 nations. Only 19.5% of trending videos reached more than one nation, but the most widely reaching, a video of a baby reacting to its mother singing (Leroux 2013), trended in 55 of the 57 nations we considered. Similarly, only 34.5% of trending videos appeared on more than one day, but some videos have a very long lifespan, such as an episode of the Senegalese telenovela *Dinama Nekh* (Senepople.com 2013), which trended for 31 days. Figure 1 shows the observed data along with the maximum-likelihood Zeta distribution (White, Enquist, and Green 2008).

We ask whether high attention towards a video within a single nation translates into increased global attention, i.e. is the home lifespan correlated to the global reach? (Brodersen, Scellato, and Wattenhofer 2012) found that social sharing was necessary for the global reach of a video, but the most socially shared videos were “trapped” in one region. We might expect either a positive or negative correlation, depending on which case is more prevalent in our data.

Our data shows only a very weak correlation between national popularity and international reach (Pearson $r = 0.20$, $p < 0.001$). Figure 2 shows the histogram of video reach and lifespan resembles the 2-dimensional Pareto distribution we would expect for entirely uncorrelated values.

We also consider the values of the reach and lifespan averaged over the videos in each nation. We found no significant overall correlation ($r = 0.20$, $p = 0.130$) but Figure 2 shows a non-uniform distribution that may suggest more complex behavior. To investigate further, we sort each nation’s videos into quadrants representing short/long-lived and local/global trends, split on the means for that nation. Figure 3 shows the results. The local video trends show no significant lifespan-reach correlation, but we see strong, significant correlations for global videos. Surprisingly, these correlations are stronger than those of all videos combined, as shown in Table 1. These results suggest that although there is no strong overall reach-lifespan correlation, there are classes of video trends for which a strong correlation exists. Videos in these classes tend to be global and short-lived relative to each nation’s typical values. The identification of such video classes is a potential area for further study.

National mean lifespan ratios show a strong, negative correlation ($r = -0.86$, $p < 0.001$) with mean global reach. The values for each nation are plotted in Figure 4. As nations become more globally connected, their lifespan ratio goes down, implying that a larger fraction of their trends

Reach	Lifespan	By Video		By Nation	
		r	p	r	p
local	short	-0.24	0.076	0.07	0.605
local	long	-0.24	0.074	0.07	0.626
global	short	0.38	0.004	0.79	<0.001
global	long	0.20	0.144	0.41	0.002

Table 1: Correlation of mean home lifespan and mean global reach.

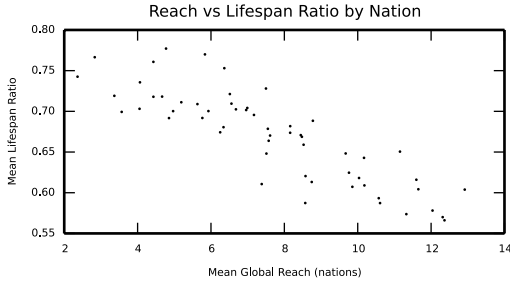


Figure 4: National mean lifespan ratios show a strong, negative correlation ($r = -0.86$, $p < 0.001$) with mean global reach. We see nations with locally-based trends that are not highly global, and nations with highly-global but externally based trends. But we see no trendsetter nations that have both locally-strong and widely shared video trends. We conclude the L.A. Effect is not dominant in YouTube trends.

begin externally. The extremes of this spectrum are nations with more locally-focused video trends (e.g. South Korea, Turkey, Morocco) and nations more likely to watch globally-popular trending videos of external origin (e.g. Canada, New Zealand, Switzerland). The L.A. effect hypothesis would suggest the existence of exporter nations which are highly globally connected (high global reach) with an inward focus (high lifespan ratio), but we do not see any examples of such nations, and conclude that the L.A. effect, if present, is not dominant.

Network Analysis

Using the symmetric conditional co-affiliation, we reduce the nation-video affiliation network into a nation-nation co-affiliation network, shown in Figure 5. We characterize the network using degree distribution, eigenvalue and betweenness centralities, and hierarchical clustering. Our analysis reveals a decentralized and highly-interconnected network. We find a subset of English-speaking nations to be highly-central hubs. We also see several Scandinavian nations with high centrality. The U.K., South Africa, and in particular, the U.A.E. emerge as key bridges between international communities.

We characterize nations in the similarity network using standard centrality measures. A high eigenvector centrality signifies highly-connected “hub” nations, while a high betweenness centrality signifies “bridges” that connect otherwise weakly-connected groups of nations. To identify bridge

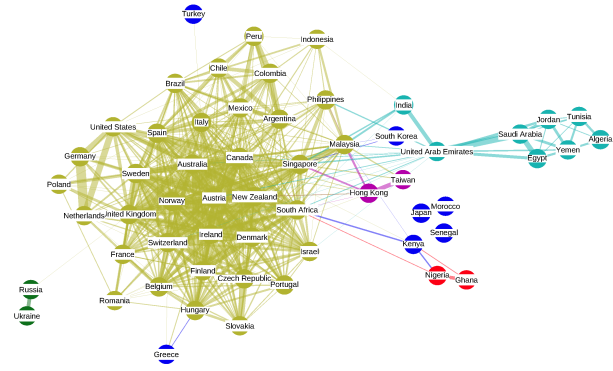


Figure 5: Edge weights correspond to symmetric conditional co-affiliation. Only edges stronger than a threshold ($-\log_2 s_{i,j} < 2.26$ bits) are included. Nodes colors correspond to communities found using single-linkage hierarchical clustering with the same threshold.

nations that may be overshadowed by other stronger bridges, we also find a “recalculated” betweenness centrality by removing the most central node, and recalculating until no shortest paths are left. We use NetworkX¹ to find all centrality values.

The nations with the largest eigenvalue centrality are shown in Table 2. Many of the most central “hub” nations are English-speaking, consistent with Ronen’s finding of English as a hub language (Ronen 2013). Scandinavian nations also appear high on the list. We also note that a nation’s mean conditional co-affiliation exhibits a strong correlation ($r = 0.9957$, $p < 0.001$) with its eigenvalue centrality, shown in Figure 6, confirming the validity of conditional co-affiliation as a meaningful quantity.

Table 3 shows all nations with nonzero betweenness centrality. Strikingly, very few nations have any betweenness centrality, suggesting a decentralized, highly-interconnected network, in which each nations weakest and strongest ties are of comparable strength. The U.A.E. is a clear exception. Its high centrality derives from Arabic-speaking nations which are much more strongly connected to each other than to the rest of the world. The U.A.E. however, has a strong connections to all nations, and thus acts as a key global bridge for Arabic-speaking nations.

We apply single linkage hierarchical clustering (Sibson 1973) to the co-affiliation network to determine its community structure. This procedure finds the most similar pair of nations, merges them into a cluster, and repeats until there is a single cluster remaining. We can divide nations into arbitrarily fine- or coarse-grained clusters by choosing a threshold level of co-affiliation, below which, clusters remain separate. We use the hierarchical clustering implementation provided by the SciPy package².

We found that a co-affiliation information threshold of 2.6 bits (chosen by inspection) divides the network into five clusters: Russia/Ukraine; Ghana/Nigeria; Hong

¹NetworkX version 1.7. <https://networkx.github.io/>

²SciPy version 0.13.2. <http://scipy.org/>

Nation	Eigenvalue Centrality
Canada	0.2148
Australia	0.2098
New Zealand	0.2093
Austria	0.2071
Denmark	0.1920
Ireland	0.1903
South Africa	0.1896
Sweden	0.1890
Switzerland	0.1881
United Kingdom	0.1870
Norway	0.1830
Finland	0.1699
Singapore	0.1694

Table 2: Highest eigenvalue centralities.

Nation	Betweenness
U.A.E.	46(46)
Egypt	2(5)
Saudi Arabia	0(4)
United Kingdom	2(2)
Netherlands	2(2)
Canada	1(1)
Sweden	0(1)
Israel	0(1)
Tunisia	0(1)

Table 3: Highest betweenness centralities of the complete network and with higher centrality nations removed (in parentheses).

Kong/Taiwan; the Middle East and North Africa; and a large cluster containing much of the Western World and Southeast Asia (see Figure 5). Several nations are not strongly associated with any cluster: Kenya, Greece, Turkey, South Korea, Japan, Morocco, and in particular, Senegal. A closer examination of our data reveals that many of the trending videos in Senegal are episodes of a popular telenovela (Senepople.com 2013), suggesting that prominent locally-produced content may play a role in Senegal’s particularly weak international co-affiliation. It is worth emphasizing that although these clusters divide along cultural lines, as described by (Huntington 1993), content is shared across clusters nearly as often as it is shared within them.

Correlates of Video Co-Affiliation

We present a multiple regression analysis of video trend co-affiliation as a function of bilateral factors: migration co-affiliation (The World Bank 2011), geographic proximity (GeoNames 2014), shared major language, shared major religion (Central Intelligence Agency 2014); and unilateral factors population (The World Bank 2014b), population density (The World Bank 2014b; Central Intelligence Agency 2014), total and per capita GDP (The World Bank 2014a), immigrant percentage (The World Bank 2011), global con-

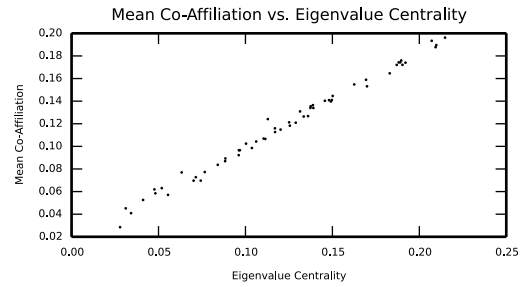


Figure 6: Mean symmetric co-affiliation shows a strong correlation with eigenvalue centrality in our data ($r = 0.9957$, $p < 0.001$).

nectedness (Ghemawat and Altman 2014), Internet penetration (Central Intelligence Agency 2014), language diversity (Lewis et al. 2014). We apply a logarithmic transformation when the data are highly non-normal. For factors representing probabilities, including co-affiliation, we apply a log transformation to the self-information of the probability, rather than to the probability itself. We calculate two models: Model 1 includes GDP per capita, while Model 2 includes total GDP. We found total GDP to be more significant for our data. We report both standardized coefficients β_i (with independent variables scaled to have a variance of 1) and unstandardized coefficients b_i that have not been rescaled. The standardized β_i show effect size relative to a standard deviation of change in the input factor. Generally, low and high variables will have different variances, so the unstandardized b_i allow comparison of effect size relative to absolute changes (e.g. an n-person change in population). The results of our analysis are summarized in Table 4.

We find that both bilateral migration and geographic proximity are significant. Bilateral migration shows one of the strongest positive correlations with co-affiliation. Perhaps surprisingly, we find that geographically distant nations tend to have higher co-affiliation. International diaspora likely play a role in connecting distant nations.

As each nation pair has two factors for each unilateral factor, and each may be positively or negatively correlated with co-affiliation, there are four possible hypotheses of how a unilateral factor can correlate with co-affiliation. These four hypotheses align well with the four international communication hypotheses presented in (Norris and Inglehart 2009).

When both the lower and higher variables in a nation pair are positively correlated with co-affiliation, increasing a factor for either nation increases their level of co-affiliation. This hypothesis corresponds to the Bangalore effect, as it suggests the factor is generally correlated with greater connectivity. We find that this is the only hypothesis consistent with both total GDP and GDP per capita. We also find language diversity and global connectedness to be consistent with the Bangalore effect.

When both variables are negatively correlated with co-affiliation, increasing a factor for either nation decreases their level of co-affiliation. This configuration suggests the factor is generally correlated with less connectivity, and cor-

		Model 1				Model 2			
		β_i	S.E.	b_i	S.E.	β_i	S.E.	b_i	S.E.
Population [†]	(low)	-0.14267***	0.02944	-0.15754***	0.03250	-0.37755***	0.05008	-0.41690***	0.05530
	(high)	-0.11662***	0.02800	-0.09890***	0.02374	-0.12504	0.06828	-0.10603	0.05790
Pop. density	(low)	-0.01640	0.02314	-1.93254	2.72743	-0.00946	0.02202	-1.11476	2.59437
	(high)	-0.18367***	0.02331	-0.94202***	0.11955	-0.20185***	0.02228	-1.03529***	0.11427
GDP per capita [†]	(low)	0.34295***	0.04619	0.38004***	0.05119	-	-	-	-
	(high)	0.06698	0.04254	0.14756	0.09371	-	-	-	-
Total GDP [†]	(low)	-	-	-	-	-0.09097	0.06242	-0.08290	0.05688
	(high)	-	-	-	-	0.41825***	0.04121	0.44512***	0.04386
Migrant % [‡]	(low)	-0.29680***	0.03144	-0.65087***	0.06895	-0.24622***	0.02980	-0.53995***	0.06535
	(high)	-0.23833***	0.03423	-0.57270***	0.08226	-0.10564**	0.03469	-0.25385**	0.08336
Connectedness	(low)	0.28031***	0.04046	2.18143***	0.31487	0.19689***	0.03903	1.53225***	0.30375
	(high)	0.04980	0.03463	0.38368	0.26677	-0.01253	0.03245	-0.09652	0.25004
Internet Pen.	(low)	0.09158*	0.04579	0.41712*	0.20856	0.24872***	0.04414	1.13285***	0.20106
	(high)	-0.14889***	0.04097	-0.85175***	0.23439	-0.13075***	0.03332	-0.74797***	0.19063
Language diversity	(low)	0.04096	0.03055	0.05640	0.04207	0.02318	0.02872	0.03193	0.03956
	(high)	0.14171***	0.02765	0.15029**	0.02933	0.12541***	0.02634	0.13301***	0.02794
Bilateral migration [‡]		0.29612***	0.03328	1.06290**	0.11945	0.29815***	0.03157	1.07020***	0.11334
Proximity [§]		-0.13341***	0.02750	-0.14535***	0.02997	-0.12705***	0.02621	-0.13843***	0.02856
Observations		1104				1104			
Adjusted R ²		0.4985				0.5445			

Table 4: Correlates of international video trend similarity ($-\log I_{j,k}$). Similarity is calculated as the negative, log-transformed self-information of co-affiliation. Each observation consists of a nation pair and includes two variables per unilateral factor (e.g. the higher and lower GDPs). The β_i are the usual standardized coefficients, calculated with all independent factors scaled to a variance of 1. The b_i are unstandardized coefficients, representing the same model, but with low/high variable pairs on the same scale for comparison. †: Transformed by taking natural log. ‡: Transformed by taking negative natural log of self-information. § Proximity is calculated by taking the natural log of the inverse distance. ***: $p < 0.001$, **: $p < 0.01$, *: $p < 0.05$, .: $p < 0.1$.

responds to the Taliban effect. We find that population is only consistent with this hypothesis, suggesting that more populous nations tend to share fewer trends with other nations. Surprisingly we also find that this is the only hypothesis consistent with the observed correlation of migrant fraction, suggesting that a higher migrant population has a negative effect on exposure to trends from other nations. We also found population density to be consistent with the Taliban effect.

When co-affiliation is positively correlated with the higher value, and negatively with the lower, a higher difference between the two nations correlates to a higher co-affiliation, corresponding to the L.A. Effect. This hypothesis can also be likened to cultural imperialism or the World Systems Theory (State et al. 2013), in which dominant cultural exporters or trendsetters overshadow the local cultures of other nations. We also find the behavior of total GDP and language diversity to be consistent with this hypothesis (among others).

In the remaining hypothesis, co-affiliation is negatively correlated with the higher value, and positively with the lower. As nations become more similar, they are more likely to share similar content, corresponding to the Firewall effect. This hypothesis is the only one consistent with the behavior of Internet penetration in our data, suggesting that nations tend to be exposed to video trends from other na-

tions with similar levels of Internet penetration. Population density, GDP per capita, and global connectedness were also consistent with this hypothesis.

Discussion

One of our most significant findings is the high level of decentralization and interconnection in the video trend co-affiliation network. Existing literature on digital communication suggests that international communication patterns cluster according to pre-existing cultural boundaries (e.g. Huntington civilizations) (State et al. 2013; Taneja and Wu 2014). While our findings are consistent, we find that the cultural boundaries are much less significant than the high level of global connectivity. We interpret this as evidence that online video enables cross-cultural communication not just in theory, but also in practice.

We also find that video trends are not being set by trendsetter or cultural exporter nations. This conclusion is supported by our analysis of video trend lifespan and global reach, with globally-connected nations always favoring videos that trend most heavily externally. While (State et al. 2013) observed exporter/trendsetter behavior in international email data, we find that GDP is most consistent with the Bangalore effect in our trending video data, corroborating the absence of L.A. effect trendsetters.

Our analysis of contextual factors shows support for all

	Population	Population Density	GDP (per cap.)	GDP (total)	Migrant %	Connectedness	Internet penetration	Language diversity
L.A.	-	-	-	✓	-	-	-	✓
Bangalore	-	-	✓	✓	-	✓	-	✓
Taliban	✓	✓	-	-	✓	-	-	-
Firewall	-	✓	✓	-	-	✓	✓	-

International communication hypotheses consistent with each unilateral contextual factor, based on the multiple regression coefficients in Table 4.

Table 5: Factors consistent with international communication hypotheses.

of the international communication hypotheses described by Norris and Inglehart. Our analysis reveals that while direct migration between two nations correlates to higher co-affiliation, high migration nations are less globally connected in general. We conclude that within our trending video data, no single model of international communication is universal. Instead, contextual factors have a multifaceted impact on the co-affiliation between different nations.

Our analysis reveals migration factors to have a surprisingly complex interaction with co-affiliation. As expected, we see higher co-affiliation between nations with high migration between them. However, the unilateral migration data reveal a more nuanced relationship. Nations with a high percentage of immigrants tend to have lower co-affiliations. In other words, migration between two nations tends to connect them while at the same time disconnecting them from other nations.

Our development of a conditional co-affiliation index, based on Ronen’s exposure index (Ronen 2013), allowed us to analyze the video trend network in several novel ways. The probabilistic nature of conditional co-affiliation gives meaning to path lengths in a co-affiliation network, allowing the use of betweenness centrality to identify bridges and quantify overall connectedness. Similarly, as a measure of similarity between nodes, the conditional co-affiliation can be correlated to bilateral data. An analysis of bilateral data may reveal patterns not visible in unilateral data, as we see with migration. As a general co-affiliation index, we believe conditional co-affiliation can be applied to better understanding many types of affiliation networks.

In general, we see that most nations are sharing diverse sets of videos with a large number of other nations, influenced by multiple competing contextual factors. While we see English-language hubs and cultural bridges like the U.A.E., U.K, and South Africa, online videos regularly cross all cultural boundaries. As nations grow in size and connectivity, their cultural interactions do not fall into a single pattern, but rather show a range of both isolating and diversifying effects, with some factors consistent with each of the L.A., Bangalore, Taliban, and Firewall hypotheses. Populous nations appear more isolated, though high-GDP nations are more more connected. Pairs of nations with high bilateral migration share more trends with each other, but fewer with other nations.

Our results raise several questions that could be addressed by future work. How is co-affiliation changing over time? As YouTube and the Internet become more mature commu-

nities, is global connectivity increasing or decreasing? Do other video and social networks show similarly high co-affiliation? And do other networks show similar correlations with contextual factors? Answering these questions will help determine whether our observations are platform-specific, medium-specific, or more general.

Conclusion

We see that in many ways, online video is achieving its potential to communicate across geographic and cultural boundaries. We also see evidence that this communication need not necessarily result in cultural imperialism or the global export of dominant media cultures. Our analysis suggests that all contextual factors cannot be explained by a single, universal model of international communication. This work contributes to a developing understanding of the factors influencing international, digital communication, which by informing design and policy, can help preserve valuable local cultures while enabling the promise of digital cosmopolitanism.

Acknowledgements

The work in this paper was carried out at the MIT Center for Civic Media, which is funded by the John S. and James L. Knight Foundation, the Open Society Foundations, the Ford Foundation, and the MIT Media Lab. We thank J. Nathan Matias, Brian Keegan, Cesar A. Hidalgo, Karthik Dinakar, Pablo Navarro, and Brandon Willard for feedback and helpful discussions.

References

- Bhattacharyya, A. 1943. On a measure of divergence between two statistical populations defined by their probability distribution. *Bulletin of the Calcutta Mathematical Society* 35:99–109.
- Bonacich, P. 1972. Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology* 2:112–120.
- Borgatti, S. P., and Halgin, D. S. 2011. *Analyzing Affiliation Networks*. SAGE Publications. 417–433.
- Breiger, R. L. 1974. The duality of persons and groups. *Social Forces* 53(2):181–190.
- Brodersen, A.; Scellato, S.; and Wattenhofer, M. 2012. Youtube around the world: Geographic popularity of videos.

- In *Proceedings of the 21st international conference on World Wide Web*. ACM.
- Central Intelligence Agency. 2014. The world factbook. <https://www.cia.gov/library/publications/the-world-factbook/>.
- Cha, M.; Kwak, H.; Rodriguez, P.; Ahn, Y.-Y.; and Moon, S. 2009. Analyzing the video popularity characteristics of large-scale user generated content systems. *IEEE/ACM Transactions on Networking* 17.
- Chang, T., and Lee, J. 1992. Factors affecting gatekeepers' selection of foreign news: A national survey of newspaper editors. *Journalism Quarterly* 69:554–561.
- Cheng, X.; Dale, C.; and Liu, J. 2008. Statistics and social network of youtube videos. In *Workshop on Quality of Service*. IEEE.
- GeoNames. 2014. Geonames. <http://geonames.org/>.
- Ghemawat, P., and Altman, S. A. 2014. DHL Global Connectedness Index. http://www.dhl.com/content/dam/Campaigns/gci2014/downloads/dhl_gci_2014_study_high.pdf.
- Gill, P.; Arlitt, M.; Li, Z.; and Mahanti, A. 2007. Youtube traffic characterization: A view from the edge. In *Proc. of the 7th ACM SIGCOMM conference on Internet Measurement*, 15–28. IMC.
- Gillespie, T. 2014. The relevance of algorithms. In *Media Technologies*. MIT Press.
- Huntington, S. 1993. The clash of civilizations. *Foreign Affairs* 72(3):22.
- Leroux, A. 2013. Emotional baby! too cute! https://www.youtube.com/watch?v=nIsCs9_-LP8.
- Lewis, Paul, M.; Simmons, G. F.; and Fennig, C. D. 2014. Ethnologue: Languages of the world. <http://www.ethnologue.com/>.
- Norris, P., and Inglehart, R. 2009. *Cosmopolitan Communications*. Cambridge University Press.
- Ronen, S. 2013. The Structure and Implications of the Global Language Network. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Senepeople.com. 2013. Dinama nekh episode 31. <https://www.youtube.com/watch?v=2vODw2maok8>.
- Sibson, R. 1973. Slink: an optimally efficient algorithm for the single-link cluster method. *The Computer Journal* 16:30–34.
- State, B.; Park, P.; Weber, I.; Mejova, Y.; and Macy, M. 2013. The mesh of civilizations and international email flows. In *WebSci 2013 Proceedings*. ACM.
- Susarla, A.; Oh, J.-H.; and Tan, Y. 2012. Social networks and the diffusion of user-generated content: Evidence from youtube. *Information Systems Research* 23:23–41.
- Taneja, H., and Wu, A. X. 2014. How does the great firewall of china affect online user behavior? isolated 'internets' as culturally defined markets on the www. *The Information Society* 30:297–309.
- Tanimoto, T. 1957. An elementary mathematical theory of classification and prediction. Technical report, IBM.
- The World Bank. 2011. Bilateral migration and remittances. <http://go.worldbank.org/JITC7NYTT0>.
- The World Bank. 2014a. GDP (current US\$). <http://data.worldbank.org/indicator/NY.GDP.MKTP.CD>.
- The World Bank. 2014b. Population, total. <http://data.worldbank.org/indicator/SP.POP.TOTL>.
- White, E. P.; Enquist, B. J.; and Green, J. L. 2008. On estimating the exponent of power-law frequency distributions. *Ecology* 89:905–912.
- YouTube. 2015. About youtube trends. <http://youtube-trends.blogspot.com/p/about-youtube-trends.html>.
- Zuckerman, E. 2013. *Rewire: Digital Cosmopolitans in the Age of Connection*. W.W. Norton & Company, Inc.