This article was downloaded by: [189.122.57.117]
On: 13 April 2014, At: 10:38
Publisher: Routledge
Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK

Porn Studies
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/rprn20

Deep tags: toward a quantitative analysis of online pornography

Antoine Mazières[ad], Mathieu Trachman[b], Jean-Philippe Cointet[a], Baptiste Coulmont[c] & Christophe Prieur[d]

[a] Institut Francilien Recherche, Innovation et Société, University Paris-Est Marne-la-Vallée, France

[b] Institut National d'Etudes Démographiques, France

[c] Centre de Recherches Sociologiques et Politiques de Paris, Université Paris 8, France

[d] Laboratoire d'Informatique Algorithmique: Fondements et Applications, Université Paris Diderot, France

Published online: 21 Mar 2014.

To cite this article: Antoine Mazières, Mathieu Trachman, Jean-Philippe Cointet, Baptiste Coulmont & Christophe Prieur (2014) Deep tags: toward a quantitative analysis of online pornography, Porn Studies, 1:1-2, 80-95, DOI: 10.1080/23268743.2014.888214

To link to this article: http://dx.doi.org/10.1080/23268743.2014.888214

PLEASE SCROLL DOWN FOR ARTICLE

Taylor & Francis makes every effort to ensure the accuracy of all the information (the "Content") contained in the publications on our platform. However, Taylor & Francis, our agents, and our licensors make no representations or warranties whatsoever as to the accuracy, completeness, or suitability for any purpose of the Content. Any opinions and views expressed in this publication are the opinions and views of the authors, and are not the views of or endorsed by Taylor & Francis. The accuracy of the Content should not be relied upon and should be independently verified with primary sources of information. Taylor and Francis shall not be liable for any losses, actions, claims, proceedings, demands, costs, expenses, damages, and other liabilities whatsoever or howsoever caused arising directly or indirectly in connection with, in relation to or arising out of the use of the Content.

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden. Terms &

Conditions of access and use can be found at http://www.tandfonline.com/page/terms-and-conditions

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Porn Studies, 2014
Vol. 1, Nos. 1–2, 80–95, http://dx.doi.org/10.1080/23268743.2014.888214

Routledge
Taylor & Francis Group

Deep tags: toward a quantitative analysis of online pornography

Antoine Mazières[a,d]*, Mathieu Trachman[b], Jean-Philippe Cointet[a], Baptiste Coulmont[c] and Christophe Prieur[d]

[a]*Institut Francilien Recherche, Innovation et Société, University Paris-Est Marne-la-Vallée, France;* [b]*Institut National d'Etudes Démographiques, France;* [c]*Centre de Recherches Sociologiques et Politiques de Paris, Université Paris 8, France;* [d]*Laboratoire d'Informatique Algorithmique: Fondements et Applications, Université Paris Diderot, France*

(*Received 28 August 2013; accepted 18 January 2014*)

Downloaded by [189.122.57.117] at 10:38 13 April 2014

The development of the web has increased the diversity of pornographic content, and at the same time the rise of online platforms has initiated a new trend of quantitative research that makes possible the analysis of data on an unprecedented scale. This paper explores the application of a quantitative approach to publicly available data collected from pornographic websites. Several analyses are applied to these digital traces with a focus on keywords describing videos and their underlying categorization systems. The analysis of a large network of tags shows that the accumulation of categories does not separate scripts from each other, but instead draws a multitude of significant paths between fuzzy categories. The datasets and tools we describe have been made publicly available for further study.

Keywords: online pornography; computational social sciences; sexual categories; network analysis

Introduction

> The purpose of these keywords rests not upon their descriptive powers, but in the potential of naming. Naming creates both the symbology and the actuality of the world. (Sigel 2000, 12)

When Linda Williams compared different kinds of pornography, revealing a proliferation of 'diff'rent strokes for diff'rent folks' (1992), she shed light on both historical and political phenomena. Indeed, during the 1970s there was a shift from a dominant male audience for pornography (Kendrick 1987) to diversified publics, along with the appropriation and staging of new desires. This ongoing diversification has been a central aspect of contemporary pornography, although it has been relatively unexplored.

Recently, this trend has been further amplified in line with a more general diversification of information sources and content, fostered largely by the development and democratization of the web and of media editing tools (Shirky 2008; Weinberger 2007). These have opened up niches for producers and broadcasters

*Corresponding author. Email: antoine.mazieres@gmail.com

© 2014 Taylor & Francis

targeting a wide range of specific sexual desires (Williams 2004). The development of user-generated content has also contributed to the blurring of boundaries between amateur and professional, mainstream and alternative, and has permitted a variety of fantasies to be showcased (Jacobs 2007; Paasonen 2010).

However, this proliferation has not been accompanied by a study of its dynamics. In Williams' early article, sadomasochistic, homosexual and bisexual pornographies are taken to illustrate the gap between the norm and 'perversity', without taking into account the new interactions between categories that stem from their co-existence. It is the specificity of niches rather than the relations between them that is explored; for example, the appearance of new fantasies and their social background (Williams 2004), or the development of alternative pornographies (Jacobs, Janssen, and Pasquinelli 2007; Taormino et al. 2013). But online pornography triggers new questions and internet activity provides logs of users' activity, allowing quantitative analysis on an unprecedented scale. Traces left by billions of users give us cultural snapshots of tastes and, more importantly, they enable researchers to look for structures and patterns in the evolutionary dynamics of practices adopted by a significant and growing proportion of the human population. As Hendler et al. note: 'A large-scale system may have emergent properties not predictable by analyzing micro technical and/or social effects' (2008, 2). This opens the way for a 'computational social science' (Lazer et al. 2009), drawing on skills from various disciplines for processing computations on huge corpuses and interpreting their results with accuracy. This approach has been applied to many fields of inquiry, such as language dynamics (Lieberman et al. 2007), evolution of science (Chavalarias and Cointet 2013), culture (Michel et al. 2011), social networks (Easley and Kleinberg 2010), and epidemic forecasting (Ginsberg et al. 2008).

The availability of data from online platforms makes pornography a good candidate for such an approach. By collecting data on thousands of videos from two main pornographic platforms, we collected a large dataset of pornographic keywords and the relationships between them (where links exist between keywords that have been applied to the same videos). Our study focuses on categorization rather than consumption practices (Attwood 2005; Bozon 2012; Wright 2013), porn production (Edelman 2009; Trachmann 2013) or the images themselves. The fact that the keywords are not randomly distributed means that they represent elementary atoms of information. If we were to postulate that 'words inform sexuality' (Sigel 2000), our research explores the possibility that 'porn tags inform pornography'.

Our hypothesis is that classification is not an organization of separated and hierarchical categories, as a Durkheimian perspective would suggest (Durkheim and Mauss 1901). It is not reducible to a virtuous circle, with practices and categories reinforcing each other and certifying the 'good' sexuality of those who are only heterosexual, monogamous, vanilla, and so on, as described by Rubin (2011). However, it does not follow that classification is anomic. In our datasets, discrete categories are related to each other and the whole system of relations exhibits a 'fuzzy logic'. The accumulation of categories does not separate fantasies from each other, but permits flow from one fantasy to another and draws thousands of paths corresponding to more and more precise desires. The proliferation of pornographic categories not only adds minor fantasies to major fantasies; it also shows how hegemonic desires provide a path to other desires, and how these other desires can be subsumed in hegemonic ones.

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Several studies have applied quantitative schemes to traces from online pornography. Amanda Spink et al. (2004) analyzed the logs of two former web search engines for the year 2001 and identified the frequency of sexual queries within the whole corpus of web search, along with the most frequent terms associated with them. The proportion of specific queries for illegal pornography, such as child pornography in peer-to-peer networks, has also been studied (Latapy, Magnien, and Fournier 2013). In addition, general case studies with weblogs from several networks have been presented with collateral analysis of porn use. For instance, Berker (2002) analyzes a German university network and makes some observations about the volume and characteristics of porn-related traffic with respect to the network as a whole. A similar, more extended application of this approach can be found in the work of Ogas and Gaddam (2012), who analyzed 400 million search-engine queries in order to unveil the 'billion wicked thoughts' of its users.

In this article we present the methods used to acquire our datasets and their main characteristics, and go on to focus on the underlying classification systems and the structural differences they imply. Online content categorization has been the focus of many studies of online interaction and collaboration (Guy and Tonkin 2006; Cattuto et al. 2009). We recall one of their major structural elements, namely the highly skewed distribution of the categories: a large proportion of items are covered by a very small number of almost universal categories, while a long tail of more specific categories still gather a large variety of content (Anderson 2006). This phenomenon encourages great diversity in content and induces the development of niches (Brynjolfsson, Hu, and Smith 2006). We explore various methods for analyzing categories, ranging from frequency measurement to network analysis, in order to reveal the diversity behind hegemonic categories, and the means by which the interactions within them are assembled into niches.

Classifying one's desire: dataset acquisition and description

Online porn is available in numerous forms. Because of their small size, plain text stories, picture galleries and comics were probably the first types of porn content to be widely diffused on the web. Audio and video files came later, with video the main medium during the 2000s, largely due to the wider availability of broadband internet connections and better streaming technologies that have enabled us to view, upload and host videos easily. However, video-hosting platforms are in competition with other kinds of services (Ogas and Gaddam 2012) that enable direct interaction between pornographic actors and viewers. For example, LiveJasmin.com, a webcam-based interaction platform, is ranked as the third most visited website in the adult category.[1] Webcam communities broadcast unstructured content – often streamed video and chat – which is unarchived and has little metadata. Despite the importance of this growing medium of online pornography, the lack of structure in the data means it is outside the scope of our study. Video-hosting platforms, on the other hand, present well-structured data. Every video belongs to a page, with a specific URL, a list of associated keywords and various other metadata such as the number of views, upload date, comments, votes, descriptions, and so on. This information is publicly available to any user, and the method we used to collect our data differs from that used by a regular user only in its systematic approach.

Downloaded by [189.122.57.117] at 10:38 13 April 2014

According to several website popularity rankings,[2] we identified the two most popular pornographic video hosting platforms – XNXX and XHamster. We created a dedicated computer program to carry out the navigation and data collection tasks required to gather the metadata for all available videos on both websites without downloading any videos. The datasets are available online[3] and are released under a Creative Commons License.[4] As shown in Tables 1 and 2, a variety of data is attached to each entry. The last column indicates how much of the dataset's entries are provided with the data described in each row.

The XNXX and Xvideos[5] domains are the oldest among the most popular porn platforms, dating from 1997. In July 2013 the websites claimed to host more than 3.5 million videos. We gathered information for 1,166,278 videos that were uploaded before March 2013. XNXX releases very little data about the videos it hosts. As shown in Table 1, only the title, keywords and comments are available to the public. Information about uploaders and the number of views is hidden or not logged by the platform maintainers. Our interest in this dataset lies primarily in its tags. When someone uploads a video, they can attach any number of keywords to their file. These keywords are meant to describe the video and highlight its specificities in order to help the user find it more easily, by anticipating the words used in a search query targeting this content. By allowing uploaders to index their videos with numerous keywords, XNXX possesses a corpus of over 70,000 tags. Among the most common pornographic platforms, XNXX is the only one to have such a corpus of descriptive keywords.

XHamster is a recent platform dating from 2007, and probably for this reason hosts fewer videos. All of the videos can be accessed, and our dataset includes all of the videos hosted by the platform since its creation and still available when we collected the data in February 2013. This represents 786,121 entries in the format described in Table 2. The presence of a timestamp on 99% of the videos permits analyses of changes through time.[6] To avoid taking incomplete years into account while considering metadata evolution, years 2007 and 2013 are omitted. An anonymized identifier links the uploader to their video clips. This permits us to track the repetition of videos among uploaders and the relations between uploaders with specific content categories or video characteristics (e.g. runtime, comments, views).[7]

As two of the most important pornographic platforms, XNXX and XHamster offer a representative sample for studying online pornography. Moreover, the structure of their data is significantly different, which makes them amenable to a comparative approach.

Table 1. Description of XNXX dataset.

Data ID	Description	Percentage of the dataset
title	Title of the video	100
nb_comments	Number of comments posted on this video	99
tags	List of the keywords associated with this video	93

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Table 2. Description of XHamster dataset.

Data ID	Description	Percentage of the dataset
title	Title of this video	100
upload_date	Day when the video was uploaded	99
channels	List of the keywords associated with this video	99
Nb_views	Number of times this video has been displayed	99
Nb_votes	Number of users who voted on this video	99
runtime	Length of the video in seconds	99
uploader	Anonymized identifier of the uploader's username	95
nb_comments	Number of comments posted on this video	92
description	Description attached to this video	48

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Categorization systems

Tags, categories and keywords are similar words for semantic descriptors. They are fundamental elements of the contemporary web: they sort content into menus and lists. They are the basis of the algorithms that allow content to be indexed in such a way as to improve the searching and browsing experiences of users. On pornographic platforms, keywords may describe practices ('BDSM', 'blowjob'), ethnic or cultural characteristics of actors (nationalities, geographical region, skin colour, religion), places (bus, bedroom, public places), devices (bed, dildo), filming techniques ('point of view', 'hidden', 'hd') and so on (Tan Hoang 2004; Attwood 2010). The keywords define the degree of semantic diversity available to uploaders in their content descriptions, and to viewers in their search queries.

On both XNXX and XHamster, videos are categorized by their uploaders. However, the platforms have different categorization systems. XHamster has a traditional top-down system that limits uploaders to pre-determined categories for characterizing their content, and viewers correspondingly only have these categories available for identifying content. This is the most common approach to categorization in pornographic platforms, most of them providing a similar list of 'classic' categories. XNXX has a bottom-up approach, letting uploaders choose their own words to index their videos, resulting in a list of more than 70,000 so-called 'tags'. This system offers greater semantic variety to the viewers, facilitating the emergence of keywords and their combinations.

The difference between top-down categories and bottom-up tags is characteristic of changes in classification strategies and practices in the digital era (Bowker and Leigh Star 1999; Weinberger 2007). The latter – known as 'folksonomy' – is a key feature in the development of content diversity and, in our case, in the tracking of contemporary porn diversification (Attwood 2007). The substantial difference in the range of semantic possibilities for uploaders and viewers impacts the number of dimensions indexed by the platforms and is therefore observable in our study.

However, despite the two platforms having different categorization systems, there are some strong similarities between the datasets, which suggests a possible generalization to other pornographic platforms. One structural similarity is that whatever the number of categories available, a very small number of tags allows one to access most of the content. For instance, on XNXX the top 5% of the most

popular tags covers more than 90% of the videos. On both XHamster and XNXX the most frequent categories, respectively 'amateur' and 'blowjobs', target 30% of all entries. To further explore the datasets beyond the identification of the few dominant widespread categories, we designed several other methodological tools.

From frequency to network

Behind this structure lies a 'long tail' of less common sexual scripts and descriptors, calling for finer-grained approaches. We first rank tag frequencies by their occurrence in titles, or using alternative methods. Then, taking into account the highly skewed distribution of tags, we shift our focus to the relationships between them. Network analyses of these relationships allow us to monitor the dominance of certain tags, revealing the diversity of the porn semantics network and the niches within the network.

The hegemony of high frequencies

Word frequency in titles

All of the videos possess one title describing their content. Some recurring archetypes (such as 'boss', 'secretary', 'maid', 'brother's best friend', etc.) can be identified in the datasets. The words 'mom' or 'mother' are present in 37 of the 100 most seen videos in XHamster. Therefore, while our study focuses on more structured aspects such as categories, we have released a tool[8] for plotting and comparing word frequencies over time in video titles from the XHamster dataset (Figure 1).

The fact that titles are unstructured sequences of characters poses challenges for conducting a systematic analysis. Spelling and typing errors, abbreviations, uses of plural and conjugated forms can all result in significant biases. For word frequencies in XHamster's titles, our algorithm strips out dashes and catches any occurrence of the query in the title; for example, 'blow' catches 'blowing', 'blowjobs', and so on while leaving biases from typing errors ('blwjob') and abbreviations ('bj')

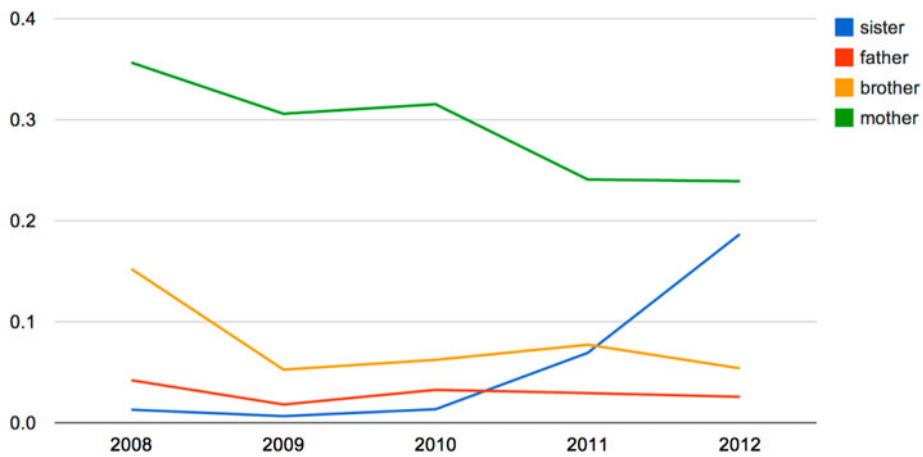

Figure 1. Example of query over Porngram.

Downloaded by [189.122.57.117] at 10:38 13 April 2014

unhandled. In this example, adding typing errors and abbreviations increases the number of blowjob videos by 16%.

Category frequencies

For tag frequencies in XNXX, our algorithm only catches the specific instance of the query, which means 'blowjob' will only catch the tag 'blowjob' (case insensitive). By considering [blowjob(s), blowing, bj, blow(s), blow-job(s), blowwjob, blwjob] as variants of 'blowjob', we increase the number of videos considered in XNXX by 5%. The bias induced by typing errors and abbreviations is thus significantly lower than for word frequencies in titles, even though our algorithm catches no variants. This phenomenon is induced by *folksonomies* (Halpin, Robu, and Shepherd 2009; Cattuto et al. 2009) where uploaders tagging their videos make a greater effort to use the most common descriptor than when they are writing titles.

We can rank categories by their frequency of occurrence; that is, for each tag, the number of videos having that tag (most videos have several tags). The top keywords represent the descriptors from which most of the videos can be accessed. If they illustrate strong practices or cultural trends, they may also overlap with other categories and get their dominant position from the transversality or generality of the concept they refer to. For example, 'amateur' and 'blowjob' do not exclude many other categories, such as those derived from sexual practices, nationalities, ethnic groups, scenarios, and so on. Adding other dimensions to the ranking by occurrences allows us to highlight interesting properties of pornographic content descriptors.

Popularity ranking is only available for XHamster and reveals categories by the number of views generated by all videos in a given category, weighted by the number of these videos. This shows the repetition of views on videos in a given category, revealing the consistency of viewers' requests for this content. These categories may point to content for which demand surpasses what is offered by uploaders.

User reaction ranking tends to increase the average number of comments per video of the given category. This uncovers viewers' reactions and interactions around the video's content. Without reading the actual comments, it is difficult to determine whether, for example, the reactions are simply descriptive or not. However, some videos may trigger comments and discussion.

Table 3 only provides the top 10 tags for each of the rankings, but we have released the dataset for all tags to permit further studies to be carried out.[9] Ranking tags allows us to isolate the various properties of specific porn content descriptors compared with the others. However, this focus tends to mute the high number of tags that, while not among the most frequent, still have significant levels of popularity in terms of number of videos. Taking tags into account, co-occurrences provide a far finer-grained tool for analysis, as detailed below.

Porn semantics as a network

Link over-representations: 'blowjob' does not make it 'funny'

The majority of videos in our dataset are attached to more than one category. If we consider the presence of several categories for the same videos as a link between each of these keywords, then we can build a global 'semantic' network. Categories are

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Table 3. Various ranking methods over tags, top 10.

Occurrences		Popularity	User reaction	
XHamster	XNXX	XHamster	XHamster	XNXX
Amateur	blowjob	Grannies	Cuckold	muslim
Men	hardcore	Old + Young	Midgets	hijab
Teens	amateur	Korean	Grannies	arabic
Hardcore	teen	Matures	Bisexuals	step
Blowjobs	cumshot	Arab	Strapon	tribadism
Anal	anal	Midgets	Cream Pie	girlontop
Big Boobs	brunette	Massage	Shemales	arabe
Masturbation	blonde	Swingers	Matures	cody
Matures	pussy	Italian	Old + Young	cumglass
Cumshots	sex	Turkish	German	sister

Downloaded by [189.122.57.117] at 10:38 13 April 2014

nodes connected through an edge (link) when two categories are significantly 'close' to one another. Such an analytical framework, known as network analysis and coming from the study of social relationships (Scott and Carrington 2011), has become very popular in many fields (Easley and Kleinberg 2010; Newman 2010).

As we have observed, tag frequencies are highly heterogeneous. This is the reason why we cannot simply rely on a raw count of co-occurrences to assess the relation strength between two tags. While we are aiming at describing only preferential relationships, very frequent tags such as 'amateur' or 'blowjobs' would obviously co-occur with any other tag. A measure of proximity must be defined for capturing how much the actual number of co-occurrences deviates from the theoretical value one would expect with no correlation between tags.[10] By doing so, we focus on edges between strongly connected tags.

As an illustration, 'midgets' – a low-frequency category in XHamster – is present 10 times more than expected in all videos having the tag 'funny'. This indicates a strong relation between these two categories and tells us that it is highly likely that midgets appear mainly to fulfil a 'funny' aspect of the scene. The fact that 'midgets' appears more with 'blowjobs' than with 'funny' is statistically expected and therefore ignored, while the relation between 'midgets' and 'funny' is unexpected and consequently highlighted in the network.

Given this methodology, we can look at link over-representation for each category without dominant categories swamping awareness of the strong and meaningful symbolic associations between less frequent categories.[11] Taking into account link over-representation reveals widely adopted symbolic associations between categories of the considered pornographic content.

These strong relations might illustrate obvious associations, such as tools or practices for a given behaviour, geographical region or ethnicity for a nationality, and so on. They allow more surprising observations when types of categories are mixed; for instance a nationality with an object or a practice. To illustrate such associations, we took the administrative and political entity named by categories (which we considered to be the common chunk of cultural entities) and identified their privileged relations with other types of categories. Table 4 shows the three

Table 4. Example of link over-representation between categories (XHamster).

Nationality category	Three most over-represented associated categories
Japanese	Asian, Massage, Bukkake
German	Vintage, Gothic, Grannies
French	Arab, Anal, Gangbang
British	Stockings, Bukkake, Celebrities
Russian	Babysitters, Old + Young, Teens
Indian	Arab, Asian, Emo
Brazilian	Latin, Anal, Black and Ebony
Italian	Celebrities, Vintage, Old + Young
Turkish	Arab, Funny, Celebrities
Czech	Spanking, POV, Old + Young
Thai	Asian, Massage, Squirting
Korean	Asian, Chinese, Hidden Cams
Chinese	Asian, Korean, Japanese
Swedish	Danish, Vintage, Gothic

Downloaded by [189.122.57.117] at 10:38 13 April 2014

strongest links for all categories referring to a nationality. A video uploaded with a nationality category does not necessarily take place in the related country or show actors coming from it. It does not accurately inform us of a country's sexual practices, but rather serves as an indicator of how this nationality is staged in a pornographic context. These examples may be applied to the whole set of relationships between the categories to obtain more generalized, global conclusions.

Porn semantic network

Figure 2 helps visualize the whole network obtained from the XHamster dataset. Only edges whose strengths are above a given threshold have been represented. An algorithm has automatically determined this threshold such that the final network is as sparse as possible but still composed of one unique connected component. We applied a community detection method, often referred to as the Louvain algorithm (Blondel et al. 2008), to identify cohesive subsets of tags in the corpus. These 'clusters' gather densely connected tags that are relatively disconnected from the rest of the network and may form semantically coherent units. In Figure 2 each node is coloured according to the clusters to which it belongs.

As well as the statistical measures available for network analysis, one can also sketch qualitative observations from visualization to characterize the network's structure and the relations between the nodes. Some clusters are highly thematic, referring to age ('milfs', 'teens', 'matures', 'grannies'), practices such as bondage and discipline, sadism and masochism ('latex', 'spanking', 'facesitting'), context ('beach', 'voyeur', 'flashing', 'public') or nationalities ('Thai', 'Chinese', 'Korean', 'Asian'). Other clusters are more heterogeneous and mix different types of keywords, such as 'blowjobs', 'black', 'ebony' and 'threesome'. The presence of hubs between several clusters is another remarkable property, such as 'massage' or 'Danish' having links with many others clusters, strong enough to appear in this visualization.

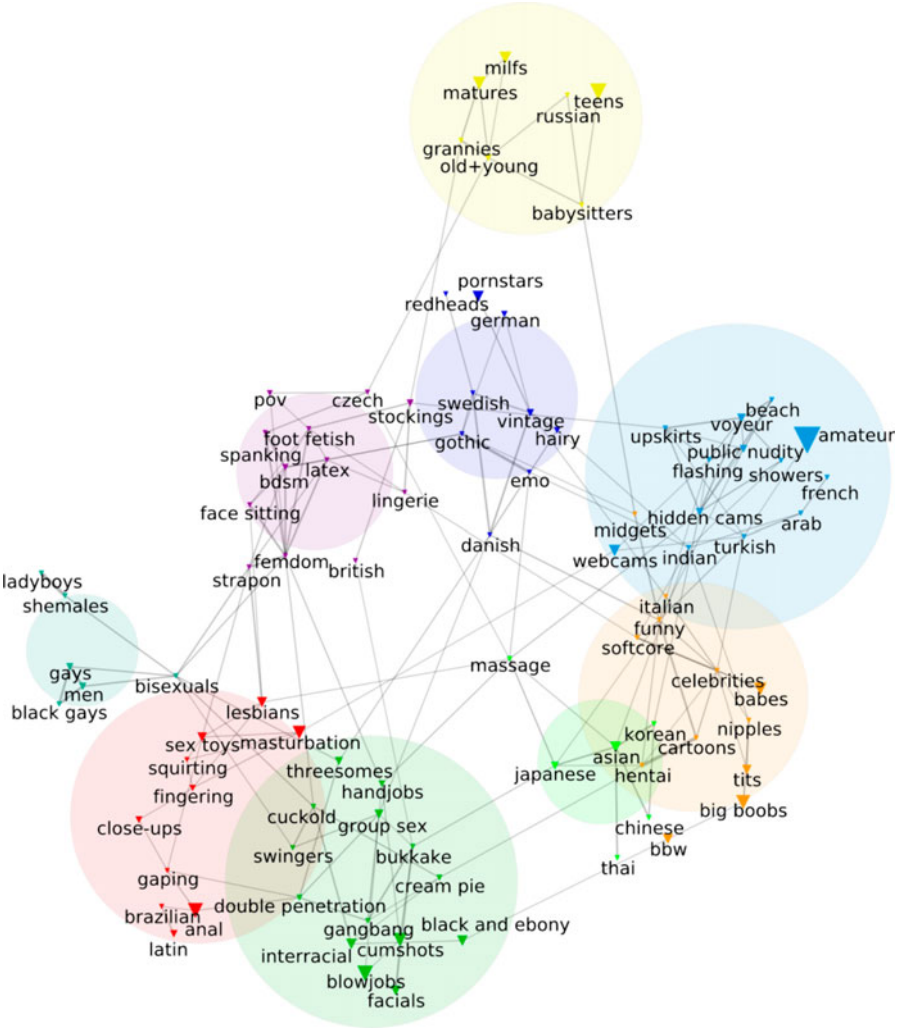

Figure 2. XHamster tag network.

Among many other possible assertions, it is worth noting the strong separation of the cluster containing the tags ‘gay’ and ‘transsexual’ from all other parts of the network. Indeed, it is connected to the rest of the network only through the tag ‘bisexual’, which constitutes a privileged bridge for any other co-occurrence. The position of the gay cluster strongly reinforces a division between heterosexuality and homosexuality by isolating the latter (Sedgwick 1990). Halperin (1995, 44) states that ‘Heterosexuality defines itself without problematizing itself, it elevates itself as a privileged and unmarked term’, so what is ‘not heterosexual’ must be defined. It therefore acquires more semantic influence upon the repertoire of desires and fantasies available on pornographic platforms. This isolation of ‘gays’ calls for a more general analysis of cases where some categories or groups of categories become to some degree peripheral to the network and constitute niches.

Downloaded by [189.122.57.117] at 10:38 13 April 2014

On Category Nicheness and Dataset Limits

We observed on the previous network that some nodes have high degrees (i.e. many links) and occupy relatively central positions in the network, while others are only connected to a few other tags and seem more peripheral in the general picture. To measure such a property more rigorously we designed a so-called nicheness coefficient. The nicheness coefficient is built upon the global matrix of mutual information between pairs of tags. We simply define the nicheness score of a tag as the sum of the preferential links connecting this tag to its relevant neighbours. The rationale behind such a measure is that tags with a 'niche' behaviour – that is, tags compatible with only few other tags – will be connected by very strong edges. Conversely, tags that may be used in conjunction with any other tags are likely to have many weakly connected neighbours and a degree of distribution that is close to random, thus resulting in a very low nicheness score. Put differently, the nicheness score also measures how much the probability of using a tag is dependent or not on the presence of other tags. If this probability remains largely unchanged with different tag pairings, the tag nicheness score is low. If the presence/absence of another tag strongly increases/decreases (and *vice versa*) the probability to observe a tag, then the tag has a higher nicheness score.

Figure 3 shows a scatter plot of the 92 XHamster channels according to frequency and nicheness. The label size scales with tag degree and node colours are consistent with Figure 1. We observe that 'hentai' and 'cartoons', although compatible with a respectable number of tags, still have a very 'biased' distribution of co-occurrences, leading to one of the highest nicheness scores. Similarly, 'ladyboys' and 'shemales' feature high nicheness score but have very low degrees (namely one and three). It is interesting to note that niche tags are not necessarily rare. 'Men' is among the 10 tags with the highest nicheness score and is the second most frequent channel. A higher nicheness score corresponds to tags that target more specialized resources. In contrast, low nicheness score tags are compatible with many other tags, and therefore provide less certain and/or less fine-grained descriptions of the content.

This empirical measure of nicheness improves upon the usual descriptions of porn niches. The niches described in Williams (1992) are practices such as bondage and discipline, sadism and masochism that are situated outside Rubin's virtuous circle (2011: 152) and practices akin to perversions of vanilla sex, whereas the many niches of online porn are in a state of flux and stem from the mobilization of specialized resources. It is not shifts in which perversions are put on/scene that form the basis of this specialization of niches, but rather specialization within major and minor sexual practices and identities (Penley 2004).

Online pornography consumers are unlikely to be immobile in the landscape of niches described by Figure 3. Some niches bring users to other niches; some of them might even attract newcomers, while others might repel viewers from porn. The paths of users within the search space should exhibit patterns relevant to understanding their 'careers' as porn consumers. Structured computer traces and other data from hundreds of millions of consumers would provide material to study pornography on an unprecedented scale. However, due to the fact that the traces left by users (mainly identification and geolocalization) on the platforms' servers are possessed by the owners of the hosting sites and are not publicly available, our dataset does not include data directly linked to users' behaviours. Access to such

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Downloaded by [189.122.57.117] at 10:38 13 April 2014

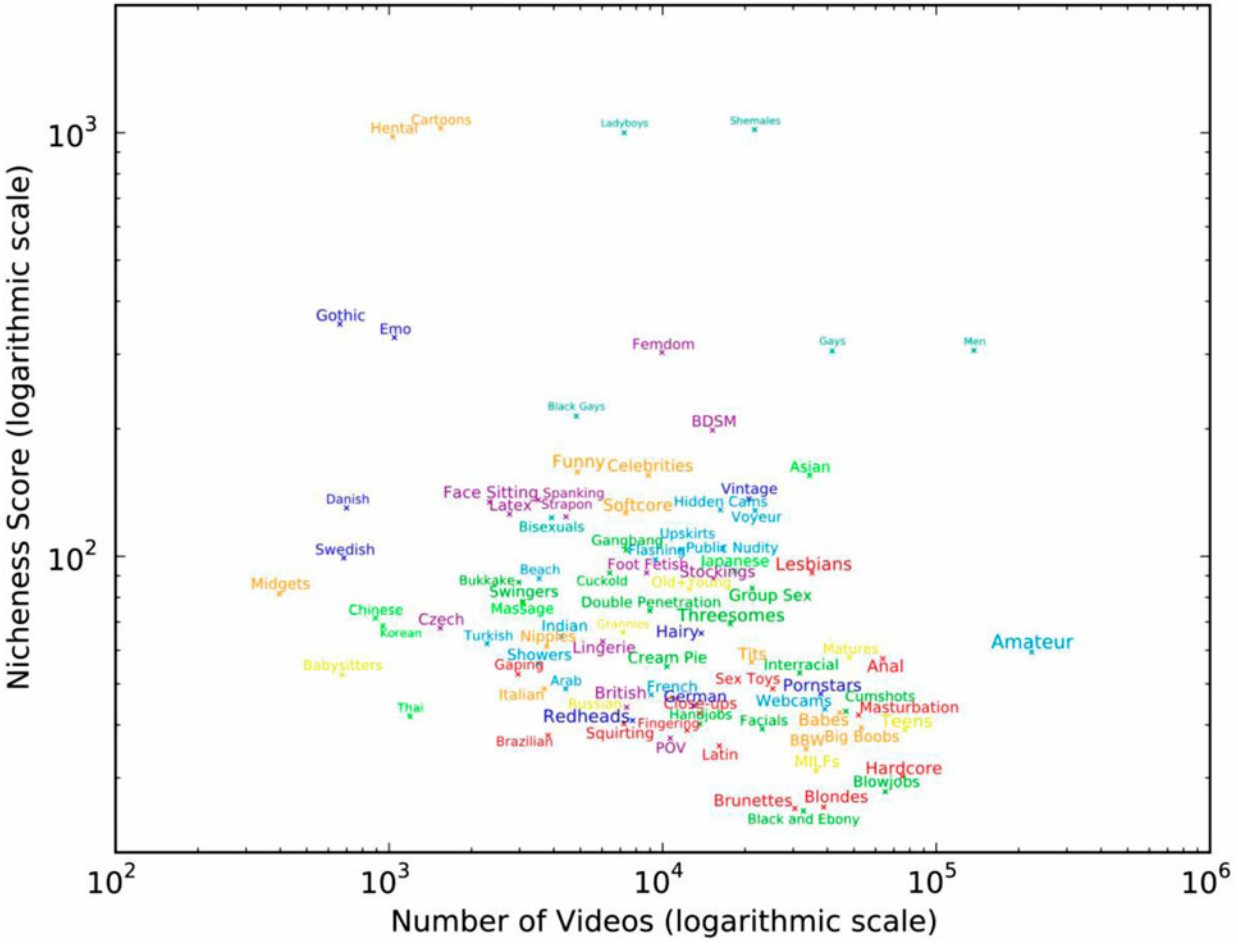

Figure 3. Nicheness of XHamster categories.

data would extend our approach and shed light on the symbols linking niches through first-hand observation of users' careers within this content.

Furthermore, tags can have different meanings in different contexts. Uses of porn categories greatly depend on national and geographical context. For example, the 'Beurette' (Arab girl in French) category is not understandable in isolation from an understanding of the French colonial past and postcolonial contemporary relationships, which produce young Arab girls as objects of desire for a white male gaze (Fassin and Trachman 2013). The potential nicheness of 'Beurette' in France could be compared with the mainstreamness of 'Arab' in North Africa or Middle East regions. We could say the same thing for the apparently most transparent 'gay', whose application varies with the different meanings of heterosexual/homosexual binarism and with cultural contexts of moral, law and sexuality. Accessing geolocalized information would therefore help to contextualize different semantic elements within their cultural surroundings.

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Conclusion

By focusing on publicly available data, this study has sought to determine whether porn tags provide a way of informing research on pornography. Such an approach does appear to help us shed light on the structural properties of porn tags so as to identify the widespread presence of dominant categories and to reveal diversity in the 'long tail' of less common sexual scripts. Beyond this general view of porn semantics, we analyzed its more discrete descriptors, involving specific users and their privileged interactions with other words. These words and their specific layouts yield heterogeneous communities of practices, objects, actors and places that inform pornography.

Our goal, using a massively quantitative approach to these phenomena, was not only to measure dominant versus under-represented categories, but to look at categorization practices in pornography. By modelling and visualizing these data, we enabled qualitative assessments to be made of tags' positions in networks and the links between categories, and therefore of how practices, nationalities, places and techniques are staged in the pornographic landscape. Large datasets and tools permit more statistical explorations and validation, but also allow a qualitative approach to be taken with respect to their numerical and visual outputs. A small-scale approach to large-scale results is likely to provide richer and more detailed information on specific communities and users.

Our study reverse-engineers users' 'tastes and colours' through the analysis of platform structures and uploaders' behaviours. While highly relevant for both website maintainers and content diffusers in devising strategies to target users, users' practices are not well understood because their traces are owned and kept by the websites. However, platform maintainers have carried out several initiatives.[12] Beyond the obvious 'buzz' and 'safe for work' marketing strategies, whose purpose is to encourage people to discover and discuss the existence of such and such platform, the data and related analyses are not verifiable. But these leaked user traces serve as evidence confirming the existence of these data in the hands of platform maintainers and their unexplored scientific potential. Allowing researchers to access these data would allow a wide range of possibilities for understanding how pornography is used and the aspects of human sexuality it represents.

Our interdisciplinary study presents the initial results of more long-term research that aims to articulate the possible contribution of large-scale quantitative methods to the theoretical and analytical frameworks provided by porn studies to understand pornographic contexts and actors. By making our datasets, analysis and tools publicly available, we hope to make this approach more accessible to those wishing to extend this approach and/or to focus more specifically on particular communities and practices, or on other aspects of porn.

Notes

1. http://www.alexa.com/topsites/category/Top/Adult. Accessed August 27, 2013.
2. Alexa and Netcraft rankings, accessed in August 2013.
3. http://pornstudies.sexualitics.org/#datasets. Accessed August 28, 2013.
4. https://creativecommons.org/licenses/by/3.0/deed.en_US. Accessed August 28, 2013.
5. XNXX and Xvideos are two interfaces to the same corpus of videos.
6. For instance, the average runtime has been multiplied by seven. Also, runtime varies a lot between categories (23 minutes for 'double penetration' and four minutes for 'men').
7. Our dataset covers the contributions of 90,000 uploaders; one-half of them being one-time uploaders only, representing only 10% of the videos.
8. http://porngram.sexualitics.org/. Accessed August 28, 2013.
9. http://pornstudies.sexualitics.org/#catrank. Accessed August 28, 2013.
10. More precisely, denoting $n(i)$ as the number of videos featuring tag i and $n(j)$ as the number of videos in which j is mentioned. The edge strength is defined as the ratio between observed and theoretical values of videos using both i and j, which can be computed as $s(i,j) = [n(i,j)N] / [n(i)n(j)]$, where N is the total number of videos.
11. The full dataset is available online: http://pornstudies.sexualitics.org/#link. Accessed August 28, 2013.
12. PornMD released an interface to explore the 10 most queried tags by country: http://www.pornmd.com/sex-seach. Pornhub, since June 2013, regularly release data and exploration tools on their data: http://www.pornhub.com/insights/. TorrentFreak looked at porn queries coming from specific countries: http://torrentfreak.com/priests-watch-dvd-screeners-while-pirates-download-filth-in-the-vatican-130407/. All sites accessed August 28, 2013.

References

Anderson Chris. 2006. *The Long Tail: Why The Future of Business is Selling Less of More*. New York: Hyperion.

Attwood, Feona. 2005. 'What do People with Porn? Qualitative Research into the Consumption, Use and Experience of Pornography and Other Sexually Explicit Media.' *Sexuality & Culture* 9 (2): 65–86.

Attwood, Feona. 2007. 'No Money Shot? Commerce, Pornography and New Sex Taste Cultures.' *Sexualities* 10 (4): 441–456.

Attwood, Feona, ed. 2010. *porn.com. Making Sense of Online Pornography*. New York: Peter Lang.

Berker, Thomas. 2002. 'World Wide Web Use at a German University – Computers, Sex, and Imported Names: Results of a Log File Analysis.' In *Online Social Sciences*, edited by Bernad Batinic, Ulf-Dietrich Reips, and Michael Bosnjak, 365–382. Göttingen: Hogrefe.

Blondel, Vincent D., Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. 'Fast Unfolding of Communities in Large Networks.' *Journal of Statistical Mechanics: Theory and Experiment* 10: 10008.

Bowker, Geoffrey, and Susan Leigh Star. 1999. *Sorting Things Out: Classification and its Consequences*. Boston, MA: MIT Press.

Bozon, Michel. 2012. 'Sexual Encounters and Sexual Practices: A Widening Repertoire.' In *Sexuality in France. Practices, Gender & Health*, edited by Nathalie Bajos and Michel Bozon, 243–264. Oxford: The Bardwell Press.

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Brynjolfsson, Erik, Jeffrey Hu Yu, and Michael D. Smith. 2006. 'From Niches to Riches: The Anatomy of the Long Tail.' *Sloan Management Review* 47 (4): 67–71.
Cattuto, Ciro, Alain Barrat, Andrea Baldassarri, Gregory Schehr, and Vittorio Loreto. 2009. 'Collective Dynamics of Social Annotation'. *Proceedings of the National Academy of Sciences of the United States of America* 106 (26): 10511–10515.
Chavalarias, David, and Jean Philippe Cointet. 2013. Phylomemetic Patterns in Science Evolution – The Rise and Fall of Scientific Fields. *Plos One* 8 (2): e54847.
Durkheim, Emile, and Marcel Mauss. 1901. 'De quelques formes primitives de classification: contribution à l'étude des représentations collectives.' *L'année Sociologique (1896/1897–1924/1925)* 6: 1–72.
Easley, David, and Jon Kleinberg. 2010. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge: Cambridge University Press.
Edelman, Benjamin. 2009. 'Red Light States: Who Buys Online Adult Entertainment?' *Journal of Economic Perspectives* 23 (1): 209–220.
Fassin Eric, and Mathieu Trachman. 2013. 'Voiler les beurettes pour les dévoiler: les doubles jeux d'un fantasme pornographique blanc.' *Modern & Contemporary France* 21 (2): 199–217.
Ginsberg, Jeremy, Matthew H. Mohebbi, Rajan S. Patel, Lynnette Brammer, Mark S. Smolinski, and Larry Brilliant. 2008. 'Detecting Influenza Epidemics Using Search Engine Query Data.' *Nature* 457 (7232): 1012–1014.
Guy, Marieke, and Emma Tonkin. 2006. Tidying up tags. *D-lib Magazine* 12 (1): 1082–9873.
Halperin, David M. 1995. *Saint Foucault. Toward a Gay Hagiography*. New York: Oxford University Press.
Halpin, Harry, Valentin Robu, and Hana Shepherd. 2009. 'Emergence of Consensus and Shared Vocabularies in Collaborative Tagging Systems.' *ACM Transactions on the Web* 3 (4): 1–34.
Hendler, James, Nigel Shadbolt, Wendy Hall, Tim Berners-Lee, and Daniel Weitzner. 2008. 'Web Science: An Interdisciplinary Approach to Understanding The Web.' *Communications of the ACM* 51 (7): 60–69.
Jacobs, Katrien. 2007. *Netporn: DIY Web Culture and Sexual Politics*. New York: Rowman & Littlefield.
Jacobs Katrien, Marije Janssen, and Matteo Pasquinelli, eds. 2007. *Click Me: A Netporn Studies Reader*. Amsterdam: Institute Of Network Cultures.
Kendrick, Walter. 1987. *The Secret Museum: Pornography in Modern Culture*. Berkeley: University of California Press.
Latapy, Matthieu, Clemence Magnien, and Raphael Fournier. 2013. 'Quantifying Paedophile Activity in a Large P2P System.' *Information Processing & Management* 49 (1): 248–263.
Lazer, David, et al. 2009. 'Life in the Network: The Coming Age of Computational Social Science.' *Science* 323 (5915): 721–723.
Lieberman, Erez, Jean-Baptiste Michel, Joe Jackson, Tina Tang, and Martin A. Nowak. 2007. 'Quantifying the Evolutionary Dynamics of Language.' *Nature* 449 (7163): 713–716.
Michel, Jean-Baptiste, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K. Gray, Joseph P. Pickett, and Erez Lieberman Aiden. 2011. 'Quantitative Analysis of Culture Using Millions of Digitized Books.' *Science* 331 (6014): 176–182.
Newman, Mark. 2010. *Networks: An Introduction*. Oxford: Oxford University Press.
Ogas, Ogi, and Sai Gaddam. 2012. *A Billion Wicked Thoughts: What The Internet Tells Us about Sexual Relationships*. New York: Penguin.
Paasonen, Susanna. 2010. 'Labors of Love: Netporn, Web 2.0 and the Meaning of Amateurism.' *New Media & Society* 12 (8): 1297–1312.
Penley, Constance. 2004. 'Crackers and Whackers. The White Trashing of Porn.' In *Porn Studies*, edited by Linda Williams, 309–320. Durham, NC: Duke University Press.
Rubin, Gayle. 2011. *Deviations. A Gayle Rubin Reader*. Durham, NC: Duke University Press.
Scott, John, and Peter J. Carrington, eds. 2011. *The Sage Handbook of Social Network Analysis*. London: Sage.
Sedgwick, Eve Kosofsky. 1990. *Epistemology of the Closet*. Berkeley: University of California Press.

Downloaded by [189.122.57.117] at 10:38 13 April 2014

Shirky, Clay. 2008. *Here Comes Everybody: The Power of Organizing without Organizations*. New York: Penguin.

Sigel, Lisa. Z. 2000. 'Name Your Pleasure: The Transformation of Sexual Language in Nineteenth-century British Pornography.' *Journal of the History of Sexuality* 9 (4): 395–419.

Spink, Amanda, Andrew Koricich, Bernard J. Jansen, and Charles Cole. 2004. 'Sexual Information Seeking on Web Search Engines.' *Cyberpsychology & Behavior* 7 (1): 65–72.

Tan Hoang, Nguyen. 2004. 'The Resurrection of Brandon Lee: The Making of a Gay Asian American Porn Star.' In *Porn Studies*, edited by Linda Williams, 223–270. Durham, NC: Duke University Press.

Taormino, Tristan, Celine Parreñas Shimizu, Constance Penley, and Mirielle Miller-Young, eds. 2013. *The Feminist Porn Book. The Politics of Producing Pleasure*. New York: The Feminist Press.

Trachmann, Mathieu. 2013. *Le travail pornographique*: *enquête sur la production de fantasmes*. Paris: La Découverte.

Weinberger, David. 2007. *Everything is Miscellaneous: The Power of the New Digital Disorder*. New York: Henry Holt Company.

Williams, Linda. 1992. 'Pornographies On/Scene, or Diff'rent Strokes for Diff'rent Folks.' In *Sex Exposed: Sexuality and the Pornography Debate*, edited by Lynne Segal and Mary McIntosh, 233–265. London: Virago.

Williams, Linda. 2004. 'Porn Studies: Proliferating Pornographies On/scene: An Introduction.' In *Porn Studies*, edited by Linda. Williams, 1–23. Durham, NC: Duke University Press.

Wright, Paul J. 2013. 'U.S. Males and Pornography, 1973–2010: Consumption, Predictors, Correlates.' *Journal of Sex Research* 50 (1): 60–71.

Downloaded by [189.122.57.117] at 10:38 13 April 2014