

# Everything is a Recommendation Netflix, Altgenres and the Construction of Taste

Emily Lawrence

Graduate School of Library and Information Science, University of Illinois at Urbana-Champaign,  
501 E. Daniel Street, Champaign, IL 61820, U.S.A., <elawrnc2@illinois.edu>

Emily Lawrence completed a Master of Library Science at the University of Maryland, College Park, and is a current doctoral student at the Graduate School of Library and Information Science (GSLIS) at the University of Illinois, Urbana-Champaign. Emily is the 2013 winner of the Miriam Braverman Prize for the article “Loud Hands in the Library: Neurodiversity in LIS Theory & Practice.”



Lawrence, Emily. **Everything is a Recommendation: Netflix, Altgenres and the Construction of Taste.** *Knowledge Organization*. 42(5), 358-364. 24 references.

**Abstract:** Netflix is a popular subscription streaming service for film and television content. The system uses non-canonical genre classifications called “altgenres” to recommend clusters of works to individual users based on their explicit and implicit preferences. Although altgenre recommendations tend to appear maximally scientific and value-neutral, they are in fact interpretive arguments in service of corporate goals. This paper explores some rhetorical elements and potential ethical ramifications of altgenres *qua* arguments.

Received: 16 July 2015; Accepted: 16 July 2015

Keywords: Netflix, altgenres, user preferences, movies, recommendations, aesthetic

## 1.0 Introduction

In her essay on believability in collections, Melanie Feinberg (2012, 333) states that, “while scholars of knowledge organization have long contended that information artifacts [like nonnarrative information systems] present arguments, such research has not emphasized how this occurs.” This paper is intended to serve partly as a corrective measure in this regard, extending Feinberg’s work on classification as a rhetorical space to a critical discussion of Netflix genre classifications. I maintain that these classifications—which Netflix has termed “altgenres” internally—possess a special kind of rhetorical power, forming a critical persuasive component of the personalized collections that Netflix presents to users.

Altgenres function as arguments about and for particular taste preferences and aesthetic experiences. They are particularly compelling because they appear scientific and value-neutral: recommended altgenres emerge from algorithmic processes and are based on a plethora of data related to users’ actual viewing behaviors, rating practices,

and other interactions with the system. However, coming to see altgenres as interpretive arguments rather than objective products of computation problematizes our initial intuitions about their reliability. Altgenre recommendations are not, then, sound reflections of some stable set of aesthetic preferences; rather, as arguments they discursively mold our taste even as our taste molds them. I contend that this ongoing process has fundamentally ethical consequences. These are largely bound up in altgenres’ status as rhetorical devices designed to maximize user engagement and retention. Understood as such, they may contribute to the gradual degradation of users’ aesthetic autonomy and the cultivation of habits of passive consumption that restrict aesthetic pleasure.

## 2.0 What are Altgenres?

For the purposes of this paper, altgenres can be understood as composite, non-canonical genres constructed from Netflix’s controlled vocabulary, with membership determined via a vast store of proprietary metadata about

individual works. There is a great degree of diversity apparent in the range of altgenres. Some examples include: gory b-horror movies from the 1980s; gay & lesbian political dramas; detective movies from the 1970s; critically-acclaimed dark French-language movies; and showbiz dramas based on real life.

Journalist Alexis Madrigal (2014) studied the descriptive infrastructure of Netflix in order to build an altgenre generator of his own. In the process of doing so, he discovered that there are 76,897 altgenres and that, further, these “are just the surface manifestation of [a] deeper database.” The “deeper database” to which Madrigal refers is the result of meticulous human labor. Netflix employs individuals with some film or entertainment-related expertise to tag every show and movie available through its services using terms (“microtags”) from its vocabulary (Fritz 2012). These employees are also trained to rate “movies [and shows] on their sexually suggestive content, goriness, romance levels, and even narrative elements like plot conclusiveness” (Madrigal 2014). Only a subset of the resultant metadata is made explicit to Netflix users in altgenre labels. Much of it is utilized in less overt ways, such as ranking within categories (Amatriain and Basilico 2012).

Madrigal found that altgenres follow a basic formula: (Region + Adjectives + Noun Genre + Based On... + Set In... + From the... + About... + For Age X to Y + Miscellaneous). Of course, individual altgenres only include some of these descriptors. Upon a fairly brief examination, we can see straight-forwardly how this formula maps onto any particular altgenre:

Gritty Military Action & Adventure from the 1980s  
 {Gritty} {Military} {Action & Adventure} {from  
 the 1980s}  
 Adjective + Adjective + Noun Genre + From the...

The altgenre formula combines more typical subject headings with “appeal elements.” In Joyce Saricks’s foundational work on Readers’ Advisory, appeal elements are terms used to describe the “feel” of a work or genre (Saricks 2005, 41; Saricks 2009). These elements refer to certain abstract characteristics related to pacing, characterization, story line, or frame (Saricks 2005, 43). For instance, terms such as gritty, dense, laconic, suspenseful, or evocative could all be used as appeal elements under the right conditions.

Altgenres operationalize customized vocabularies of appeal, and are thus effective generators of viewing appeal. Viewing appeal is analogous to the concept of reading appeal; meaning it is, *mutatis mutandis*, “the power to invoke interest in reading and to set off an action of reading” (Dali 2014, 24). Altgenres are just one way in which Netflix uses recommendations to produce appeal,

but they have significant import for aesthetic taste. I want to shift now to how altgenres function specifically as arguments for certain preferences and experiences.

### 3.0 Altgenres as Arguments

Netflix has two basic methods of determining users’ preferences: by asking what they prefer or by inferring what they prefer from (patterns in) discrete interactions with the system. The former is a measure of explicit or stated preferences while the latter relies on proxy measures for implicit or revealed ones. Because we cannot look unobstructed into the minds of others, implicit preferences can only ever be inferred from behavior. Netflix has largely shifted towards emphasizing implicit preferences in making its recommendations, and it extrapolates these preferences from “recent plays, ratings, and other interactions” with Netflix (Amatriain and Basilico 2012). Inference is thus a significant iterative step in Netflix’s recommendation process.

Because Netflix identifies appropriate altgenres largely on the basis of the user’s revealed—rather than stated—preferences, there is a sense in which the service seems to tell the user things about their personal taste that might not otherwise be accessible to them. This is one way in which Netflix might be said to aid users’ understanding of their own taste: on the basis of the user’s actual viewing patterns, the service illustrates for them certain features to which they typically respond positively. This leads Madrigal to call altgenres “a tool for introspection,” and once led a friend of mine to say that they were sometimes “too much truth for one night.”

However, questioning whether Netflix simply guides or invasively co-determines users’ taste complicates this seemingly positive introspective effect. It is critical to understand that altgenres are recommendations and that recommendations are themselves arguments. We can formulate these recommendation-arguments as follows (Wright 2012, 352):

- A. One makes a “diagnosis” that offers an explanation of some data (e.g., the user views a lot of action movies because they like films with violent action sequences);
- B. One forms a “prediction” based on that “diagnosis” (e.g., the user will enjoy movies with violent action sequences);
- C. One appeals to that “prediction” in conjunction with some “normative claim(s)” to generate a “recommendation” about what one ought to do (e.g., the user will enjoy movies with violent action sequences / watching movies with features

you like is good / the user should watch these movies with violent action sequences).

These arguments shape our experiences of works by picking out certain features and presenting them as justifications for a recommendation. As recommended categories, altgenres thus set our expectations of the works that they include by making certain features salient prior to viewing and obscuring others. They may also drastically alter how we apprehend the work-in-question at a somewhat more foundational level. In his landmark paper, “Categories of Art,” Kendall Walton (1970, 343) posits, among other things, that we perceive works of art in categories and, further, that how we categorize various works actually alters what features we take those works to have. In a sense, the category itself sets our evaluative standard for its member works, and this has profound implications for the effects art has on us and for our own aesthetic judgments.

So while altgenres shape our expectations and experiences of what we watch, they also provide scaffolding for our ensuing evaluative efforts. This occurs in the category-dependent manner that Walton details (where the category in which we perceive the work helps determine what aesthetic properties we attribute to it and thus affects our appraisal), but also in another way peculiar to recommendation. Altgenres are categories that purport not only to tell us something about their member works, but also something about ourselves. That is, they describe works of art just as they describe the user’s taste. In binding categories to our own individual preferences, altgenres acquire a special kind of explanatory power: once a user has watched a work in a recommended category, the altgenre begins to function as a suggested explanation for users’ enjoyment or lack thereof. Altgenres can tell their story about any pleasure or displeasure they experience while viewing a particular work. Netflix essentially mediates the common struggle to articulate why you like what you like by providing a brief account of what features you enjoy.

Confirmation bias may well play an essential role here in shoring up the user’s sense that Netflix has successfully identified their actual preferences rather than made an argument inferring certain preferences. For example, if I watch and enjoy *Dog Day Afternoon* as a token of the altgenre “Visually-striking Crime Movies from the 1970s,” I am apt to assume that Netflix got it right: visually-striking ‘70s crime movies are in fact the kind of thing I enjoy for the reasons Netflix presented to me at the outset. Yet, regardless of how much I enjoyed the film, Netflix’s recommendation-argument still might not be sound because its diagnostic premise (i.e., that I watch certain movies because they share these specific features) could still be inaccurate. Which is to say, I might have actually enjoyed any

number of different things about *Dog Day Afternoon*. Perhaps the rawness of the dialogue or the hopelessness of the protagonist’s downward spiral appealed to me most. It is therefore conceivable that a recommendation can actually obscure my own latent aesthetic justification.

There is also the more pernicious possibility that these recommendations are themselves self-actuating—that is, they may have the capacity to bring about in the user just those preferences that they argue the user already has. Here, however, it would be a step too far to think solely in terms of how Netflix’s recommendation algorithm shapes users when users do in fact shape the algorithm. When I say that users “shape the algorithm,” I do not mean simply that they generate data that straightforwardly affects the algorithm’s output. I want, instead, to refer to the possibility that, depending on the contours of the data users generate, the algorithm might deploy different methods of calculation. This could occur, for instance, when some trend in a user’s viewing practices triggers a new subroutine in the algorithm. At the very least, Netflix periodically adjusts the algorithm in response to patterns in user data; users can thus shape the algorithm non-mechanically by motivating intervention on the part of Netflix developers.

Any critical exploration of Netflix and its infrastructure must take the user/system transaction into account. In his critique of Facebook and its big data research, sociologist Nathan Jurgenson (2015) explains that:

Algorithms are made to capture, analyze, and re-adjust individual behavior in ways that serve particular ends. Individual choice is partly a result of how the algorithm teaches us, and the algorithm itself is dynamic code that reacts to and changes with individual choice. Neither the algorithm or individual choice can be understood without the other.

The Netflix user is then not an autonomous agent utilizing a fee-based service to meet pre-specified preferences: rather, the user enters into a mutual, dynamic transaction with the system, wherein their preferences structure and are structured by pervasive recommendation. Put simply, my personal taste affects what I see on my Netflix homepage, but the reverse is also true. Ultimately, I am a different aesthetic agent when and because I use Netflix.

#### 4.0 Ethical Consequences

There is a concerted effort on the part of Netflix developers to move towards a passive discovery model, wherein users no longer need to search for something to watch because the system provides them with suitably customized selections (Amatriain 2014). As the Netflix system de-emphasizes the practice of searching independently, it ac-

tively works to funnel users towards recommended titles. Altgenres are an especially effective way to do this, partly because their obviously customized nature highlights only one part of the user's transaction with the system: altgenres illustrate to users that their own explicit and implicit preferences determine what they see. That those preferences are themselves molded by what they see, and that they only see a curated reflection of a small subset of their inferred preferences, remains largely invisible to them. This model also carries the risk that users will become habituated to (and thus complacent with respect of) unreflectively "consuming" content. Such passive consumption does not require us to spend as much time honing the appreciative and interpretive skills that afford us greater aesthetic pleasure in the long term.

The success of such an operation depends in no small part on the trust users place in it. Melanie Feinberg (2012), in her examination of the believability of collections, formulates the concept of synthetic ethos. Synthetic ethos refers to how "the various textual elements that form the collection...work synthetically to produce a compelling rhetorical object" (Feinberg 2012, 336). In Aristotelian rhetoric, ethos is one of three forms of persuasive appeal, the other two being pathos (persuasion by evoking emotional responses) and logos (persuasion by rational argumentation); ethos persuades by virtue of the speaker's character, their perceived credibility and reliability (Poggi 2005). Speakers can cultivate ethos through demonstrations of practical wisdom, moral integrity, and goodwill (Feinberg 2012, 331-2). The concept of synthetic ethos allows us to talk about how information systems, rather than individual speakers, persuade their audiences to take particular courses of action.

One way that Netflix cultivates synthetic ethos is through its altgenres, which constitute recommendations for what it is we "really want to watch." In their critique of Big Data research, danah boyd and Kate Crawford (2012, 663) refer to a "widespread belief that large data sets offer a higher form of intelligence and knowledge that can generate insights that were previously impossible, with the aura of truth, objectivity, and accuracy." Altgenres reap the benefits of just this belief; other demonstrations of practical wisdom are barely needed when many users accept Netflix's show of scientific accuracy at face value. Moreover, altgenres themselves are persuasive partly because they obscure the human labor involved in their production. They are presented, more or less, as the incontrovertible product of computation and a data set that is thought to be representative of the user's preferences. Disagreeing with an altgenre is then something akin to denial.

One can also detect the "aura of truth" in the origins of Netflix's present-day classification procedures. When he first developed the tagging procedures that would produce

altgenres, Netflix's Vice President of Product Innovation, Todd Yellin, described this process as one that would effectively identify the "'quanta,' the little 'packets of energy' that compose each movie" (Madrigal 2014). Jens-Erik Mai (2010, 631-2) describes classification approaches of this sort as building on "a broader folk theory of categorization," wherein:

The core defining concept of classification limits the power of the classifier to simply identifying the characteristics of things, and grouping those things that share characteristics. Classification thereby becomes a technical process that is merely occupied with the discovery of what things are and placing them in the one system given by nature.

Referring to a vocabulary of appeal as if it corresponds to "little packets of energy" in a work of art makes tagging (a fundamentally interpretive practice) sound highly technical—or magical, depending on whom you ask—and, if performed well, indisputable. Although I take it Joseph Tennis (2013, 48) is correct in conceptualizing knowledge organization work as craft and arguing for an "artistic turn in descriptive practices," this is in conflict with the altgenre project as it has thus far been conceived; as it stands, altgenres' rhetorical potency emerges from a common folk conception of classification. But there is no straightforward calculation, no perfect accumulation of descriptors, that can guarantee a particular sort of aesthetic experience or verdict, and there is certainly no metaphysical "energy" that constitutes the work of art. Netflix's customized recommendations are persuasive partly because they utilize and foster a logic of aesthetic pleasure that is wholly at odds with any prevailing philosophical conception.

Relatedly, it is important to see that, however neutral or scientific it might seem, there are particular biases and human values embedded in the language of altgenres. Indeed, it is not possible to eradicate bias from classificatory procedures or their resultant systems (Feinberg 2007). A great deal more could be written on this following a careful reading of Netflix's vocabulary itself, but I will save the bulk of such an analysis for another paper. Suffice it to say that the Netflix controlled vocabulary is certainly limited and limiting. Much valuable work has already been done by scholars of knowledge organization to interrogate the ways in which classification schemes endorse and reinforce hegemonic value systems (Olson 2001; Furner 2007; Drabinski 2013; Koford 2014). This inevitably leads to the silencing of marginalized views, experiences, and (aesthetic and non-aesthetic) judgments. For instance, Netflix taggers categorize *Nightmare on Elm Street 2: Freddy's Revenge* in the following genres: Horror

Movies, Teen Screams, Supernatural Horror Movies, Cult Horror Movies, Slasher and Serial Killer Movies, and Cult Movies. Yet the film's status as a queer horror classic—primarily among queer horror enthusiasts—is nowhere detectable in its Netflix description. This is unsurprising, given that “queer” is not an available Netflix descriptor. The closest the vocabulary comes is the obviously exclusionary “Gay & Lesbian,” which renders invisible (or perhaps simply irrelevant to appeal) other queer identities. However, even “Gay & Lesbian Horror” does not appear to be an in-use altgenre, if it exists at all. This is not a function of what Netflix has in its collection, but rather of which sorts of categories it considers to be sufficient or sufficiently broad generators of viewing appeal.

Finally, it is unclear that Netflix would categorize the film as queer horror even if such an altgenre existed. Although it is not a subtle film, reading *Nightmare on Elm Street 2* as queer horror is still contingent on certain kinds of relational knowledge (to, e.g., interpret it as an allegorical horror film), contextual knowledge (to, e.g., know how it has been discussed by queer fans), and political commitments (to, e.g., understand the genre designation and take it seriously). Adherents to Yellin's folk theory of categorization would likely fail to classify the film in this way. Depending on how a particular tagger is situated, they may not be able to discover the film's membership in this genre merely by looking.

Because it is the very nature of classification to obscure certain viewpoints (Bowker and Starr, 2000), I suspect there are numerous other such examples. There are currently no formal mechanisms in place to allow users to contest Netflix classifications, a fact that surely emerges (at least partly) from its pretense to neutrality and limits progressive improvements to the vocabulary over time.

Even a reader who is sympathetic to the aforementioned concerns might at this point wonder: why should any of these worries be peculiar to or especially pronounced for Netflix? After all, we receive recommendations frequently and from a wide range of sources. Some of these recommendations are, like those Netflix produces, aesthetic in character, in that they speak to the aesthetic content of some work of art and/or employ an aesthetic judgment as part of the central predictive argument. This is true when, for instance, a Readers' Advisory librarian recommends a novel to a patron looking for something new to read. The librarian predicts that reading the novel will bring about desirable states in the patron by virtue of certain features of the work itself. Both the librarian and Netflix rely upon (necessarily biased) vocabularies of appeal, and in either case the content of the recommendation will likely frame the recommendee's experience of the work. That is, in both instances a third

party has discerned (however successfully or unsuccessfully) some purportedly relevant features of the object of recommendation that in turn affect the recommendee's reception of that object.

Given this, should we really think that the Netflix recommendation is different in kind from the librarian's? Netflix recommendations often differ greatly from the recommendations individuals make to one another, which are not necessarily consistent with any set of norms or standards and can thus vary widely in terms of structure and purpose; they may be, for instance, under-supported, malformed, or insufficiently personalized. But when customized aesthetic recommendations emerge as manifestations of some institutional descriptive infrastructure, they seem to have a great deal more in common than not.

Nonetheless, I maintain that there are some significant differences. One essential point of divergence between Netflix recommendations and those issued by librarians is that the former serve very particular pragmatic, corporate purposes. As Felix Salmon notes in his 2014 Reuters editorial on Netflix's recommendation algorithm, “Netflix's big problem...is that it can't afford the content that its subscribers most want to watch...as a result, Netflix can't, any longer, aspire to be the service which allows you to watch the movies you want to watch.” Salmon contrasts Netflix streaming with its DVD-by-mail service: whereas the latter allows users to procure most any title they like, the former is severely constrained by legal and economic considerations. Put simply, Netflix can only make accessible works for which it has procured streaming rights, and this excludes a great many things from its overall collection. This, in turn, restricts what users see in their customized collections.

These constraints are not, in and of themselves, problematic. Salmon (2014) claims that a restricted catalog incentivizes directing users towards works that are largely aesthetically poor or middling; this probably requires further investigation, as it is not yet clear that the legality of streaming a movie or television series reliably bears on aesthetic quality. However, problems emerge if we consider the effects of these constraints on user autonomy. When the ultimate goal is profit, legal restrictions do incentivize cultivating an environment of pervasive recommendation that facilitates passive discovery and directs users away from independent searches of the catalog. So when Amatriain and Basilico (2012) trumpet that, on the Netflix user interface, “everything is a recommendation,” they elide the fact that this is ultimately to aid user retention in a world where external forces limit what Netflix can make accessible. Although external forces also affect what libraries can make accessible, there is no systematic effort to obscure this, nor are readers' advisors' recommendations formulated to aid “user retention.” Netflix recommendations,

however, facilitate viewing practices centered on convenience rather than some overriding aesthetic preference while simultaneously concealing the fact that these practices are in service of corporate goals.

Finally, Netflix cares about the user's enjoyment only instrumentally. The recommendation-argument, presented in an earlier section, begins with a diagnosis that takes the following form: "the user views a lot of X because they like Xs that include Y feature(s)." While the argumentative structure is likely to remain basically stable over time, the 'enjoyment' piece is open to revision. For the time being, guiding users towards content they will enjoy seems to be the best practice for increasing engagement and, consequently, retention. However, it is conceivable that we might come to know or it might come to be the case that something else is better at fostering engagement. That new thing could potentially be aesthetically arbitrary. For instance, if movies that featured a preponderance of the color red in their promotional materials reliably correlated with increased viewing by all or a subset of users, there would be nothing to stop Netflix from factoring this into its recommendations. One can even imagine a (slightly dystopian) scenario wherein Netflix determines that it can increase engagement most effectively by eliciting addictive viewing behaviors in users. Salmon (2014) even suggests that something like this may already be going on, as Netflix recommendations increasingly emphasize television over film content because, the argument goes, viewers are apt to binge watch whole series.

In contrast, the librarian is—or should be—unwilling to alter their recommendations in ways that are purely engagement-maximizing and seemingly irrelevant to actual enjoyment. Although librarians certainly care about engagement (and might even value enjoyment instrumentally in the sense that it cultivates in patrons a genuine motive to read more), they are bound to a robust system of core values. The librarian's overarching commitments to intellectual freedom and non-coercive service would tend to preclude recommendations based on arbitrary features or addictive properties. The same cannot be said for Netflix or, in fairness, most any enterprise for which the ultimate motive is profit.

## 5.0 Conclusion

Too little critical attention has yet been paid to the rhetorical dimensions of recommender systems, to the ways in which their recommendation-arguments frame our experiences and depend upon classifications that necessarily "[valorize] some point of view and [silence] another" (Bowker and Starr, 2000, 5). I have argued that Netflix altgenres make arguments to users regarding their preferences, what they ought to watch in light of those prefer-

ences, and why they experience what they watch as they do. In so doing, I have outlined a series of concerns related to how these arguments shape taste preferences while appearing incontrovertible.

In the end, there is no denying that Netflix streaming is enormously popular. The service has more than 60 million subscribers in about 50 countries; approximately 40 million of those subscribers are in the United States (Spangler 2015). Furthermore, in 2013 it was estimated that, together, Netflix and YouTube constitute roughly half of all peak Internet traffic in North America (Bump 2013). And yet there have been no sustained interrogations of what Netflix's recommendation practices might mean for users' taste or for their experiences of the content they stream.

Recently, Netflix has begun to transition away from hyper-granular altgenres towards categories with fewer descriptors (e.g., Witty Movies, Cult Movies, Romantic Action & Adventure). Such a shift indicates that Netflix may have altered some part or parts of its recommendation algorithm. Perhaps highlighting personalization is no longer perceived as the most effective way to promote engagement, or perhaps new ways of emphasizing customization have proven more potent—for instance, including a "Because you watched X" statement with each recommended title. Maybe the change occurred because user engagement is improved by recommending broader altgenres with more member titles. Since one should not expect Netflix to be consistently forthcoming about the reasoning behind each of its decisions and since these decisions have actual consequences for the service's many users, we are in need of attentive, rigorous analyses of Netflix, of its descriptive infrastructure and the rhetorical strategies it employs to achieve its objectives. Here I have offered a brief account of how altgenres function as arguments with ethical implications, but it seems there is much more interpretive work to be done if we want to formulate answers to a critical emerging question: How does Netflix shape us?

## References

- Amatriain, Xavier. 2014. "Recommender Systems." *Machine Learning Summer School @ CMU*. <http://www.slideshare.net/xamat/recommender-systems-machine-learning-summer-school-2014-cmu>.
- Amatriain, Xavier and Justin Basilico. 2012. "Netflix Recommendations: Beyond the Five Stars." *The Netflix Tech Blog*. <http://techblog.netflix.com/2012/04/netflix-recommendations-beyond-5-stars.html>.
- Bump, Phillip. 2013. "Half of Internet Traffic in North America is Just to Watch Netflix and YouTube." *The Wire: News from the Atlantic*. <http://www.thewire.com/technology/2013/05/netflix-youtube-traffic/65210/>.

- Bowker, Geoffrey C. and Susan Leigh Star. 2000. *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: The MIT Press.
- Boyd, danah and Kate Crawford. 2012. "Critical Questions for Big Data: Provocations for a Cultural, Technological, and Scholarly Phenomenon." *Information, Communication & Society* 15, no. 5: 662-79.
- Dali, Keren. 2014. "From Book Appeal to Reading Appeal: Redefining the Concept of Appeal in Readers' Advisory." *Library Quarterly* 84: 22-48.
- Drabinski, Emily. 2013. "Queering the Catalog: Queer Theory and the Politics of Correction." *Library Quarterly* 83: 94-111.
- Feinberg, Melanie. 2007. "Hidden Bias to Responsible Bias: An Approach to Information Systems Based on Haraway's Situated Knowledges." *Information Research* 12, no. 4. <http://InformationR.net/ir/12-4/colis/colis07.html>.
- Feinberg, Melanie. 2012. "Synthetic Ethos: The Believability of Collections at the Intersection of Classification and Curation." *The Information Society* 28: 329-39.
- Fritz, Ben. 2012. "Cadre of Film Buffs Helps Netflix Viewers Sort through the Clutter." *The Los Angeles Times*. <http://articles.latimes.com/2012/sep/03/business/la-fi-0903-ct-netflix-taggers-20120903>.
- Furner, Jonathan. 2007. "Dewey Deracialized: A Critical Race-Theoretic Perspective." *Knowledge Organization* 34: 144-68.
- Jurgenson, Nathan. 2015. "Facebook: Fair and Balanced." *Cyborgology*. <http://thesocietypages.org/cyborgology/2015/05/07/facebook-fair-and-balanced/>.
- Koford, Amelia. 2014. "How Disability Studies Scholars Interact with Subject Headings." *Cataloging & Classification Quarterly* 52, no. 4: 388-411.
- Madrigal, Alexis C. 2014. "How Netflix Reverse Engineered Hollywood." *The Atlantic*. <http://www.theatlantic.com/technology/archive/2014/01/how-netflix-reverse-engineered-hollywood/282679/>.
- Mai, Jens-Erik. 2010. "Classification in a Social World: Bias and Trust." *Journal of Documentation* 66: 627-42.
- Olson, Hope. 2001. "Power to Name: Representation in Library Catalogs." *Signs* 26: 639-68.
- Poggi, Isabella. 2005. "The Goals of Persuasion." *Pragmatics & Cognition* 13: 297-336.
- Salmon, Felix. 2014. "Netflix's Dumbed-Down Algorithms." *Reuters*. <http://blogs.reuters.com/felix-salmon/2014/01/03/netflixs-dumbed-down-algorithms/>.
- Saricks, Joyce G. 2005. *Readers' Advisory Service in the Public Library*. 3rd ed. Chicago: American Library Association.
- Saricks, Joyce G. 2009. *The Readers' Advisory Guide to Genre Fiction*. 2nd ed. Chicago: American Library Association.
- Spangler, Todd. 2015. "Netflix Adds Record 4.9 Million Subscribers in Q1." *Variety*. <http://variety.com/2015/digital/news/netflix-adds-record-4-9-million-subscribers-in-q1-1201473151/>.
- Tennis, Joseph T. 2013. "Ethos and Ideology of Knowledge Organization: Towards Precepts for an Engaged Knowledge Organization." *Knowledge Organization* 40: 42-9.
- Walton, Kendall L. 1970. "Categories of Art." *The Philosophical Review* 79: 334-67.
- Wright, Larry. 2012. *Critical Thinking: An Introduction to Analytical Reading and Reasoning*. 2nd edition. New York, NY: Oxford University Press.