

technology review

Published by MIT

BUSINESS IMPACT


CLOUD COMPUTING

Decoding Social Media

Discovering the patterns in tweets will reshape TV, ads, and politics

Stuck With Oil Sands

Q&A: Google+ Creator



Can Nokia Be Saved?

The Data Surveillance State

A Social-Media Decoder

New technology deciphers—and empowers—the millions who talk back to their televisions through the Web.

By DAVID TALBOT

From his 24th-floor corner office in midtown Manhattan, the veteran CBS research chief David Poltrack can gaze southward down the Avenue of the Americas, its sidewalks teeming. For more than four decades, it has been his job to measure people's television habits, preferences, and reactions. In large part, this has meant following the viewing habits of Nielsen panels of TV viewers and parsing the results of network surveys on their opinions. On a late September afternoon, with fall premieres under way, his desk was strewn with color-coded opinions from 3,000 Americans who had wandered into CBS's Las Vegas research outpost, Television City, at the MGM Grand Hotel and Casino, and agreed to fill out TV surveys for the chance to win a 3-D home entertainment system.

But now he's also dealing with a growing force: the masses talking back through social media. Of the approximately 300 million public comments made online worldwide every day—about two-thirds of them on Twitter—some 10 million, on average, are related to television (though daily numbers vary quite widely). “¿Que sera two and a half men si[n.] Charlie?” one viewer recently tweeted, alluding to the replacement of Charlie Sheen by Ashton Kutcher on the CBS sitcom. “The beginning of Person Of Interest is like Jack & Ben all over again,” remarked another. (A couple of weeks later, another added: “I assume CBS will keep going with


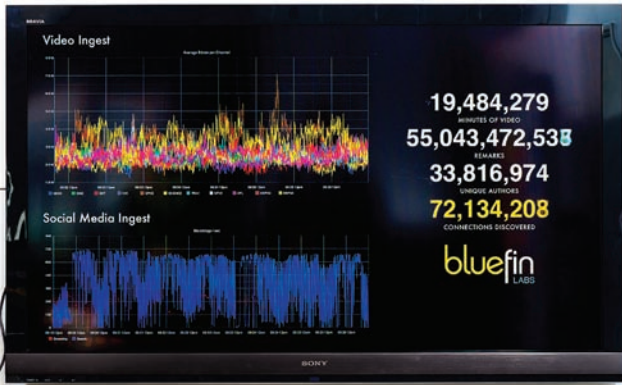
what's been working for them, and replace Andy Rooney with Ashton Kutcher.”) TV executives like Poltrack must now grapple with these spontaneous, messy, irreverent remarks.

How to make sense of it all? Poltrack walked into the office of a staff member, John Butler, clutching a report from a startup called Bluefin Labs, a social-media analytics firm that attempts to track comments on shows and ads and discern the commenters' interests and demographics. Some of what it had found seemed surprising. For example, the season premiere of *Two and a Half Men* had attracted 78,347 comments compared with 82,980 for *Dancing with the Stars*, on ABC, even though the latter show has lower Nielsen ratings and an older audience that's less likely to participate in social media. (It turns out that reality competition shows, by their nature, attract more active audience response.) Poltrack wondered how a little-watched show called *Bad Girls Club*—on the Oxygen network—had garnered 32,665 comments. “Get *Bad Girls Club* up there,” he said to Butler, motioning to Butler's computer screen. “What are they saying?” Butler scrolled through the raw comment string. “This bitch angie on #Badgirlsclub wear the same damn socks in every episode,” remarked one viewer; “BGC, shower & bed,” announced another. It was hard to know what any of it meant.

Overall, the data was raw and, in many cases, ambiguous. But Poltrack came away with some respect for what he was seeing. “As



POWER SHIFT
Deb Roy, CEO of Bluefin Labs, says social media have changed the relationship between media consumers and producers.



a one-time measurement, we have better ones,” he said, referring to CBS’s precisely constructed surveys. But whereas the surveys are intermittent, social-media analytics can provide “a continuous monitor of conversation about a program, episode by episode,” he said. “And that is something we can’t replicate.” What’s more, the quantity of commentary is increasing all the time, making it more important as an object of study and as a force network executives would like to harness. As Poltrack explained, real-world and online chatter—the “exponential movement of a conversation through the population”—drives the success or failure of TV shows and, in turn, the allocation of \$72 billion in U.S. television ad spending.

Six hundred miles to the west, a similar assessment was under way at the Cincinnati headquarters of Procter & Gamble, the world’s largest advertiser (its brands include Tide, Gillette, Bounty,

Pringles, and Duracell). Each year the company spends \$5 billion on media ads—the bulk of them on TV—and another \$5 billion on in-store advertising worldwide. While Procter & Gamble carefully vets ads with consumers before airing them, it has never known whether the same viewers would respond differently to an ad depending on what show surrounded it.

Craig Wynett, the company’s chief learning officer, says Bluefin Labs is teasing out nuances in the way context affects the extent to which an ad generates buzz. One specific product ad (he wouldn’t say which) was placed on two shows with similar demographics and ratings. One show produced eight times more social-media response than the other. Nobody knows why, but that’s what happened. “Historically, we have held context as a constant. Well, surprise! In the real world, context plays a fundamental role,” he says.



PLAYING BALL As a doctoral candidate, Michael Fleischman (above) used televised Red Sox games to teach computers to recognize home runs and other plays. Now the company he cofounded, Bluefin Labs, analyzes social media to decipher mass reactions to TV shows and ads viewed in the United States. In its offices (left), a screen displays the number of comments searched, minutes of TV ingested, and connections found.

Bluefin Labs is one of a growing number of analytics companies parsing the meaning of comments in social media. And its CEO, Deb Roy, believes they are capturing a fundamental change in the relationship between creators and consumers of mass media. “What I have learned by hanging out with TV executives, talent agencies, and creative types is that the assumption is built into their organizations’ DNA that this is a one-way dialogue,” he says. “Audience members speaking through social media is effectively a shift in power.”

In some ways, a two-way conversation has begun. And in future years a TV network could, in theory, continue the conversation by revising its promotions to emphasize characters that have caught on with audiences—or even by revising plot lines midway through a season. Advertisers, meanwhile, could swap out ads—or place them differently—on the basis of the social-media response they get. (Something like this already happens with online ads; increasingly, algorithms use real-time metrics like page views and content changes to guide placement decisions.) In the political realm, campaigns could rapidly determine, among other things, which messages animate people. And early feedback from the first adopters of analytics—network executives and advertisers—could provide clues to wider potential impacts. Wynett says he doesn’t know if the people who commented on his advertisement bought the product or “if the message spread until every man, woman, and child heard it.” Still, he says, “It’s early days, but it shows promise.”

MINING SOCIAL SENTIMENT

Analyses of online comments are already influencing corporate, financial, and governmental behavior. Certain companies, Comcast among them, keep an ear open for outbursts of anger to help them detect and respond to service outages and product problems. A London hedge fund, Derwent Capital, makes trades based on the financial calm or anxiety it gleans, in part, from social-media data. And while recent events have suggested that revolutionaries can use social media to help them overthrow some authoritarian regimes (see “*Streetbook*,” *September/October 2011*), China has learned to manage citizen outrage through measured responses to specific online complaints about matters such as police corruption (see “*China’s Internet Paradox*,” *May/June 2010*).

For marketing purposes, it has become de rigeur for companies to set up Facebook pages and send out tweets, and to keep a watchful eye on the bubbling up of blogged anger. This is true of television networks as well as other companies. For example, Discovery Communications, which runs channels including the Discovery Channel, TLC, and Animal Planet, maintains 75 Facebook pages with 45 million fans, and keeps 23 Twitter accounts crackling with reminders like “*Mythbusters starts in 5 minutes!*” “It’s all that beautiful viral effect of social media to get people to watch our shows,” says Gayle Weiswasser, Discovery’s vice president of social-media communications, “and we aren’t the only ones who do it.”

To tap the other side of the conversation—the unscripted response of consumers with social-media accounts—companies like Radian6 (now owned by Salesforce), General Sentiment, Sysomos, Converseon, and Trendrr track social-media sentiment and volume on a range of topics. Of course, even the best filtering efforts don’t eliminate all spam. And it’s not always clear what prompted a post, how a slang-filled tweet should be interpreted, or how to identify the author’s demographics. Yet it is “critically important” for businesses to make sense of all this, says Radha Subramanyam, senior vice president of media and advertising insights and analytics at Nielsen: “This is the world’s largest focus group, the world’s largest town hall. Companies that figure this out will thrive in the next 10 to 15 years. Companies that don’t will fail.”

It’s especially important for TV networks and advertisers. Nielsen says that Americans, on average, spend 20 percent of their day watching TV, and many simultaneously peck away at laptops or mobile devices. Sites like Miso and GetGlue encourage people to discuss favorite shows with friends and other fans. Evidence is emerging that social-media buzz has some relationship to ratings: NM Incite, a Nielsen-McKinsey joint venture, found that among people aged 18 to 34, a 9 percent increase in such chatter in the weeks before a show’s premiere correlated to a 1 percent ratings increase.

Recognizing these kinds of connections, sentiment-analysis firms including Trendrr.tv (part of Trendrr) and Socialguide specifically track social response to television content. But Bluefin

is unique in also tracking most of what is on TV—including the ads—to draw specific relationships between televised stimulus and social-media response. “What Bluefin is doing is technically impressive,” says Duane Varan, chief research officer at the Disney Media and Advertising Lab in Austin, Texas. Already, it’s becoming possible to measure TV viewership directly through cable boxes rather than through samples such as Nielsen panels, he says, and “Bluefin is doing a similar thing with this universe of public social-media discourse.”

THE NFL AND SOCIAL TV

Bluefin Labs’ headquarters occupy a one-story 19th-century factory that once made hoses, next to a boutique movie theater in the Kendall Square area of Cambridge, Massachusetts. Lego blocks strewn on café tables busy the fingers of visitors or employees at informal meetings. Roy, the cofounder and CEO, sits at one of an open cluster of desks in close quarters with nearly 40 employees, most of them engineers with experience in fields like artificial intelligence, search, and video analysis. A poster showing the “bloodline” of the advertising industry is pinned on a worn wooden post to his right.

Roy, who is 42, is a Winnipeg-born computer and cognitive scientist who until 2008 had spent his entire career in academia, first at the University of Waterloo and then at MIT and its Media Lab, where he became head of a research group called Cognitive Machines. Among other things, his group concerned itself with problems such as how to teach English to robots. In 2005

he launched the ambitiously named “Human Speechome Project” to document how children learn language. Before his son was born, he equipped his home with 11 video cameras and 14 microphones. Then the proud papa recorded (almost) everything that happened in the house to figure out how different adult interactions—as well as activities and objects in different locations of the house—affected the boy’s speech development. In 2008, after collecting 300 gigabytes of data every day, Roy stopped. Then he and his graduate students performed feats like charting his son’s gradual mastery of the word “water.” (A presentation of this process was the hit of the 2011 TED conference and has spread virally throughout the Internet.)


The project married linguistic analysis with video analysis, but it was Roy’s PhD student Michael Fleischman, now 34, who made the conceptual leap to TV. For his dissertation, Fleischman initially planned to use lessons from the speechome project to teach computers language. But there was a problem: “It became clear to me that I’d have to wait until Deb’s son grew up,” Fleischman says. “I needed to find a new data set.” The answer came, appropriately, in front of the television. One night, while watching a Red Sox game with his girlfriend (now his wife), Fleischman realized that televised sports had what he needed: visual action, play-by-play dialogue, and sufficient repetition and structure. So he started creating software that would turn baseball games into a language-teaching tool. The spoken words “home run,” for example—when accompanied by a camera angle arcing across a stadium—might lead the computer to learn how to distinguish an actual home run. After *Technology Review* wrote about the baseball-interpreting technology, he and Roy were invited to apply for a \$100,000 National Science Foundation Small Business Innovation Research grant.

In 2008, Fleischman and Roy got the grant and named the company after a sushi restaurant where they’d discussed their plans. The initial focus on sports led to angel investments from sports magnates including Jonathan Kraft, president of the New England Patriots; Jim Pallotta, an owner of the Boston Celtics; and Dan Gilbert, majority owner of the Cleveland Cavaliers. (As of October 2011, the company had received \$8.5 million in funding, mostly from Redpoint Ventures but also from angel investors.) Bluefin’s first customer was the National Football League, which already had a new online feature called “Game Rewind” that let fans review already-played games. Fleischman and Roy expanded the concept by tying the video stream to social-media comments. They tuned search algorithms to look for football-related keywords; the result was an on-screen interface that let fans read, play by play, what others had written. (This turned out to be an early instance of the now-popular trend in social TV applications.)

During this process, Roy and Fleischman had another “Aha!” moment. The comment stream that turned up for televised games had blank patches at regular intervals. “We looked and said, ‘What’s

CONVERSATION STARTER

The total number of social-media comments is rising sharply, providing more fodder for analysis. Most of these are public.



Source: Gnip. Figures reflect both public and private instances of active participation in online conversations: tweets, comments, and other posts.



NETWORK EFFECT David Poltrack, chief research officer for CBS, has long recognized the value of viewer conversations about shows. Now he's evaluating tools that scrutinize millions of comments made about TV online.

that?” says Roy. “Well, those were the ads.” It hadn’t occurred to them that people would talk about ads. But they do. They write—as one recently did, in a tweet picked up by Bluefin—things like *“The dude rappin in the mcdonalds commercial about the smoothies will forever be clowned where ever he goes.”*

Roy and Fleischman realized that the advertising industry might be interested in understanding more about such comments, and advertisers had large research budgets. “We took the principles of big data, data mining, and visualization,” Roy says, “and turned that microscope [in my house] into a telescope to look at the world of social media as it relates to television.” They called their work the “TV genome.” Today Bluefin has 15 clients, including Pepsi, Mars, and Comcast; the TV networks CBS, Fox Sports, A+E Networks, AMC Networks, and Turner Broadcasting; and the ad agencies McGarryBowen and Hill Holliday. The company business is selling subscriptions to its interface and custom analytics. While making these conquests, Roy encountered a language learning issue of his own. “When I started talking to people in TV, I’d hear the word ‘programming.’ Turns out they weren’t talking about programming software,” he recalls. “It took me a while to figure this out.”

INSIDE THE TELESCOPE

In order to capture almost everything happening on television, Bluefin uses a data center studded with satellite dishes in Medford, Massachusetts (see *“Heeding the Tweets,” next page*). Through the first week of October, they’d pulled in every minute of more than

210,000 episodes of 7,100 television shows, plus advertisements. The company now monitors 200 networks.

After uploading the raw feed to Amazon’s cloud computing service, Bluefin gathers programming-guide information—the names of the shows, their broadcast channels and times, and also the names of characters and actors—along with closed-captioning text extracted from the video signal itself. This provides a list of keywords that can help identify relevant social-media comments. Since advertising schedules are not published ahead of time, Bluefin creates one. The algorithm detects when a “pod” of ads has started. Then a system of digital fingerprinting identifies repeat airings; human staffers are notified of first-time airings to make the initial identification.


Among the more than 10 million comments made daily about TV content, Bluefin’s algorithms identify about 1.4 million that are made in the three hours before or after a show or advertisement aired on one of the networks it tracks. (About 90 percent of these comments are tweets; the bulk of the remainder are public Facebook posts.) Although on-demand services, recording technologies, and new Internet models of TV delivery are changing viewing habits (see *“Searching for the Future of Television,” January/February 2011*), most people still watch television the old-fashioned way, and Roy says they seem more likely to make real-time comments when they know they are watching the first airing. Bluefin also keeps close tabs on the 9.8 million people who have commented about television at least once in the past 90 days, to build up knowledge about their demographics and interests.

Text analysis underpins all these efforts: whereas “delicious” or “tasty” might indicate a positive response to a restaurant, terms like “can’t wait” or “fascinating” or “drivel” might show up in comments related to TV shows. Bluefin is working on identifying not only positive or negative reactions but ones that are vulgar or polite, serious or amused, calm or excited. “At the highest level, what we are trying to do is language understanding,” Fleischman says. It also tries to glean demographic information about who is commenting. Women, for example, are more likely to refer to family members, while men are more likely to mention friends or electronic devices. Emoticons hint at age: someone who uses :-) is probably 10 years older than someone who uses :). People using 8-) are even older.

Bluefin ultimately turns all this data into two main measurements. “Response level” reports the number of people commenting on any given ad or episode of a show, measured on a logarithmic 10-point scale. “Response share” measures what percentage of all social-media response to television programming at a given air-time focused on a particular show or ad. The company’s first interface—Bluefin Signals, which provides analytics on comments about TV shows—went live in June. A second, which is due for release in December, will track response to individual ad campaigns. Next year Bluefin plans to include Spanish-language comments in its analysis.

HEEDING THE TWEETS

How Bluefin analyzes what you say about TV



1. Video intake

Every day, Bluefin ingests video from 200 television networks, representing about 90 percent of the programming viewed by U.S. audiences. It also captures the name and time of the show, the names of characters and actors, and closed-captioning text of the show's dialogue. It tracks advertisements as well. Machines detect ads; humans electronically tag new ads, and video fingerprinting technology detects and tracks repeat airings.



2. Social-media intake

At the same time, Bluefin scans 300 million public social-media comments daily for keywords associated with the video signals it has processed. The system seeks relevant comments that appear in the three hours before or after a show is broadcast, suggesting that the words are not being used in some other con-

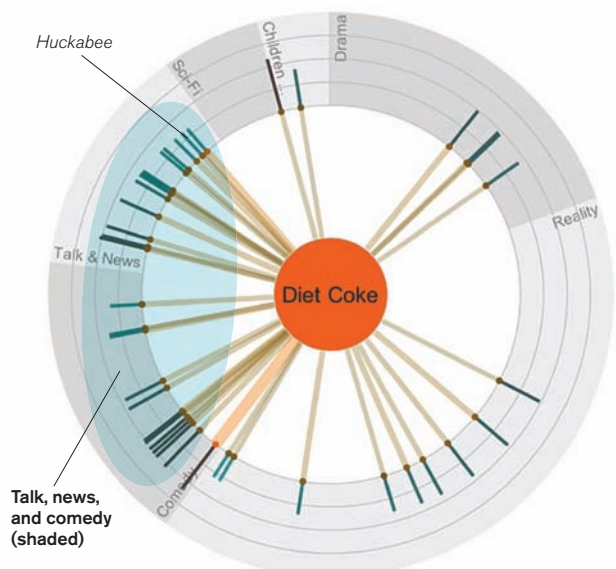
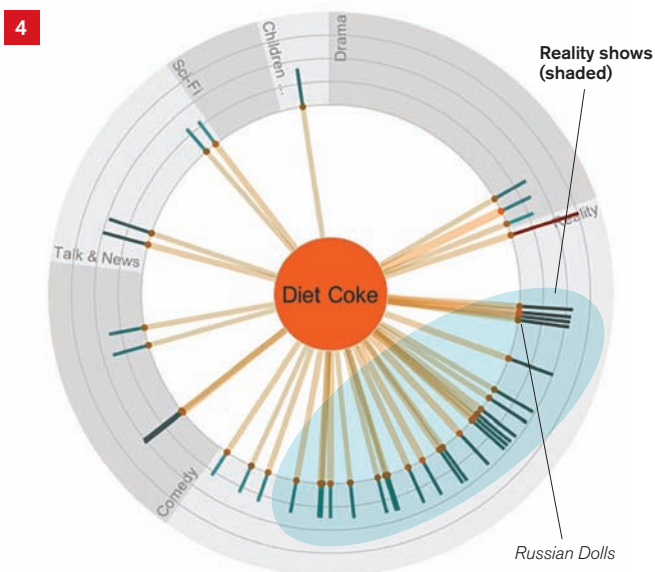
text. Each day, about 1.4 million comments fit these criteria.

3. Text analysis

For comments about TV, Bluefin seeks clues about the author's gender and age. In this example (based on a real tweet but edited and anonymized), a female screen name, use of multiple exclamation points, and references to family members are hints of female authorship. The system keeps track, in anonymized fashion, of posters' commenting habits—especially what TV shows and ads they comment on over time.

4. Patterns and associations

Bluefin makes many kinds of associations that could be valuable for programmers, marketers, and, someday, politicians. For example, women who talk about Diet Coke in social media also discuss reality shows more than other kinds of programming, with *Russian Dolls* topping the list. But men who mention Diet Coke in social media tend to discuss talk, news, or comedy shows the most, especially Mike Huckabee's program on Fox News. Such information can, in theory, do things like help ad buyers determine which slots best provoke audience "conversation," but proof of its value is still under study.



Roy says there is no reason why the company couldn't track television signals and analyze the sentiments expressed through social media in other countries, too. So far, though, there are no immediate plans to expand beyond the United States.

HOW PEPSI RESONATES

Bluefin can tell you certain things very clearly, and one of them is the degree to which audiences are moved to talk about Diet Pepsi when a swimsuit-clad Sofia Vergara is involved. Vergara, who plays the Colombian trophy wife on the comedy *Modern Family*, appeared this past summer in a widely aired commercial in which she met soccer heartthrob David Beckham on a beach. Traditional social-media analysis showed a 7 percent increase in chatter about the drink during the time the campaign aired. But Bluefin knew the ad had run exactly 746 times, on 260 different shows, and it knew who had commented on those 260 shows during the commercial's run. Among those 1.8 million people, mentions of Diet Pepsi were up by 19 percent. Bluefin was also able to determine that in June, a spike in negative sentiment about the Hyundai brand throughout social media coincided with the premiere of a TNT science fiction drama, *Falling Skies*, during which commenters complained that a promise of "limited commercials" had been broken.

Such insights could be a boon to television advertisers wondering what ads to place and where to place them. "If I'm moved by comedy, drama, or sheer creativity in an ad, then I have a propensity to talk about it," says Mike Proulx, senior vice president for social media at the ad agency Hill Holliday. "There is a theory—and it's unproven—that the greater the social-media mentions, the higher the content's resonance."

The new tools could also complement analytics like those provided by Simulmedia, a New York City company that licenses anonymized viewing data from 18 million set-top boxes. "Bluefin is able to associate specific delivery of an ad to positive sentiment in some target audience," says Dave Morgan, Simulmedia's CEO. "That alone is becoming a key marketing objective. It's no longer just 'Spend a certain amount of money on sex, age, income'; it's 'Spend money that causes positive sentiment from a target audience.'"

Of course, the applicability of Bluefin's data goes only so far. For one thing, most conversations still happen in the real world, not online; according to one market research firm, KelleyFey, 90 percent of the conversations people have about brands in the United States happen offline. Further confusing matters, "people who use social media are not representative of the general population, it's very difficult to understand the differences, and it's a dynamic, variable thing," says Varan, the Disney research executive. "There is so much we don't know about how the social-media universe differs from the real universe. So the danger is looking at the kinds of results that Bluefin would produce and drawing conclusions that it's a reflection of what the overall population is doing."


MAKING OF THE PRESIDENT, 2012

Fleischman and Roy predict, however, that applications will ultimately go well beyond TV, helping to reveal the events and media sources that inspire people in radio, newspapers, and magazines, as well as online. "You can look at the affinity from any one thing to any set of things," Fleischman says. "Pop culture expands fairly widely—politics, media, actors, books, plays, religion. The tail gets longer and longer, to anything you can imagine people talking about. If you are looking for a set of people with a particular interest, we can tell you how this relates to another set of interests."

The next obvious area for Bluefin is politics. Early next year, in time for the presidential primary season, the company expects to analyze social-media reaction to speeches, televised debates, and political advertisements. "What's potentially more interesting is understanding who, positively and negatively, is making connections between audience members," Roy says. Already, social media have become a key political organizing tool (see "*How Obama Really Did It*," *September/October 2008*). But political operatives are tough customers: they care mainly about two things. First, in presidential races, they care what undecided voters in swing states think. And second, they want to know who, among their reliable base of supporters, is willing to take some action, such as donating or spreading a message. Comments in social media are of limited value.

Still, political operatives might be just as interested as Procter & Gamble in learning which messages are resonating. Andrew Heyward, former president of CBS News and a Bluefin advisor, thinks such analyses could be vital for political commentators and anchors, too. "Getting almost real-time feedback on scale and sentiment is very valuable for a political organization or a candidate—or a news organization trying to cover the race," he says.

The value of social-media analytics will only rise. The number of comments is growing every month, and Roy predicts that analytical technology will improve as algorithms are refined and as more participants crunch the data and pursue yet-unimagined applications. Today, conversations in social media are still hard to overhear and decipher. But what might someday emerge from Bluefin, or one of its competitors, are technologies that make those conversations easy to capture and understand—and produce a metric akin to a Nielsen rating. (The Nielsen and McKinsey findings about the correlation between buzz and ratings are a step in that direction.)

In the future, then, marketing officers and network executives such as David Poltrack may be able to leave the survey takers back at the Vegas slot machines and tune in to a continuous social-media conversation that is now either inaudible or incomprehensible. They may see ways to create television programming, advertising, political communications, and ultimately other media that are smarter—or at least more responsive to what audiences find appealing. 

DAVID TALBOT IS TECHNOLOGY REVIEW'S CHIEF CORRESPONDENT.