Demand Finder: Set Top Box Television Ad Targeting using a Novel Interactive Data Visualization System

Brendan Kitts, Dyng Au, Brian Burdick, Jon Borchardt, Amanda Powter and Todd Otis[1]
PrecisionDemand
821 Second Avenue, Suite 700, Seattle WA 98104 USA
bkitts@precisiondemand.com

Abstract—**This paper will show how machine learning and data visualization techniques are being used to execute real television ad buys. We present an innovative data visualization tool which allows users to filter, histogram, and sort so as to identify the television inventory with highest value per dollar. Using the application users have been able to identify media that performs 50% better than previous campaigns as measured by phone response in several live television campaigns.**

Keywords—television; ad; targeting; set top box; visualization

I. Introduction

TELEVISION currently attracts about 65 billion dollars in ad spending per year in the United States, and is growing in consumption along with other video-related mediums. Despite its size, television ad targeting technology has remained little changed from the 1950s, and has continued to use age-gender matching as measured using the Nielsen panel, a group of 25,000 persons who document their viewing activities [1][2][7]. However in just the last few years, some significant changes have begun to take place. One of the biggest is the emergence of return path Set Top Box data. From effectively 0% of US Households in 2008, the number of Set Top Boxes with return path capability has increased to over 30% in 2012 in the United States[2] [8]. This explosion in behavioral television data means that it is possible to target with a precision that has never before been possible.

This paper will show how interactive data visualization techniques can allow Media Buyers to build highly targeted, low cost TV ad campaigns based on multi-million dollar ad spending budgets.

II. The TV Ad Targeting Problem

The problem that we want to solve is to select a set of TV media into which to insert an ad, such that advertiser value per dollar is maximized. Let M_i be a contiguous segment of time in the TV video stream that a station is offering for sale, $CPI(M_i)$ be the cost per impression of the timeslot, $r(M_i)$ be the buyers per impression and $I(M_i)$ be the impressions for the timeslot. The objective is to select a set of Media which maximizes:

$$\sum_i r(M_i) \cdot I(M_i) \qquad (1)$$

subject to $\sum_i CPI(M_i) \cdot I(M_i) \leq B$ and $V(\{M_i\}) = true$

Where B is the television campaign budget, V determines if the set of media violates rotation rules (such as running an ad more than once per 60 minutes, having greater than 5% of budget on any one network or day-part, and so on). Rotation rules are defined by television ad buyers. A simple greedy strategy for allocating television media is to select media in order of value per dollar

$$M_i: \max \frac{r(M_i)}{CPI(M_i)} \qquad (2)$$

subject to the rotation rule constraints until the budget is filled. In order to implement the above strategy, we need to use Set Top Box data to estimate $r(M_i)$. We also need to estimate the historical clearing price of the media $CPI(M_i)$.

III. Targeting Algorithm

The value of media $r(M_i)$ is defined as the number of buyers per impression whom will be predicted to be viewing a particular media segment. Unfortunately directly estimating buyers based on historical observations runs into a matching problem – the probability of an individual person being a buyer of a product and also of being detected in the viewing population is low. For example consider an advertiser with 100,000 product purchases (1 person in 3,000 per US population); and with 1 million set top boxes (1 viewer in 114 US TV households). Persons view on average 6 hours per day, so in any one day there is a 1 in 4 chance of a person being present. A large broadcast airing has 1,000,000 impressions. We could then expect 1/3,000 * 1/114 * 1/4 * 1,000,000 = 0.7 buyer-viewers to be detected on a television broadcast airing. This means hopefully 1, but often 0 buyers would be in the audience for a given program airing, which means that we would be working with very low statistics.

We solve the matching problem by mapping individuals in Set Top Box data and Buying data into two anonymized vectors in high-dimensional (3,500 element) demographic

---

[1] Kitts, B., Au, D., Burdick, B., Borchardt, J., Powter, A., Otis, T. (2013), Demand Finder: Set Top Box Television Ad Targeting using a Novel Interactive Data Visualization System, *Proceedings of the Thirteenth IEEE International Conference on Data Mining Workshops (ICDMW-2013)*, Dallas, TX, IEEE Press.

[2] As of 2012, 30% of households have set top boxes that are Tru2Way compatible, and another 30% of households have Direct Broadcast Satellite Set Top Boxes that are also believed to be return path capable [8].

space. We then calculate the degree of match of this buyer vector with media vectors [3][4][5]. The matching process is shown in Figure 1. The method has some appealing privacy protections since personally identifiable information can remain with its primary sources, and targeting fundamentally is based on anonymous, aggregated data.

The method enables more data to be used for matching because the entire set of persons contribute to the buyer demographic vector and media vector. This enables rare demographic variables (eg. “diabetes interest”, “spectator sport tennis”, etc.) to participate in the match. We use cosine similarity to estimate the degree of match

$$r(\bar{M}_i) = \frac{\bar{P}^{+} \cdot \bar{M}_i^{+}}{\left|\bar{P}^{+}\right| \cdot \left|\bar{M}_i^{+}\right|} \qquad (3)$$

where $\bar{P}^{+}$ and $\bar{M}_i^{+}$ are buyer and media vectors respectively. After calculating the degree of match, we then estimate cost of media and impressions

$$CPI(M_i) = \sum_T h_T \cdot CPI(\bar{m}_{i,T}) \qquad (4)$$

$$I(M_i) = \sum_T g_T \cdot I(\bar{m}_{i,T}) \qquad (5)$$

where h_T are coefficients which minimize squared error between predicted CPI and actual based on historical record of spot prices, and $\bar{m}_{i,T}$ comprise features of the airing, such as the CPI average for station, CPI average for time-of-day, CPI average for program name, and so on. g_T is defined similarly for historical impressions as measured by Set Top Boxes.

Figure 2, 3, 4, 5 and 6 show an example of television ad targeting for a handyman tool by advertiser Rockwell Tools Inc. Figure 2 shows a screen capture of the television ad, Figure 3 the buyer demographic profile, Figure 4 and 5 best matching media and Figure 6 best value per dollar media.

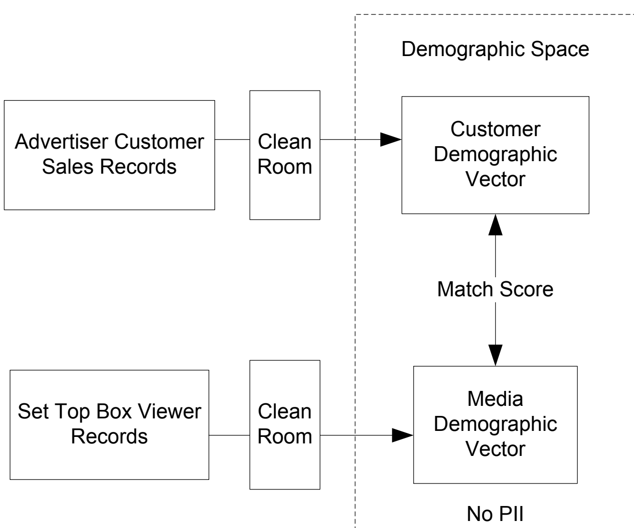

Fig. 1. Targeting Algorithm

Fig. 2. A frame from one of Rockwell Tools’ 30 second US television commercials. The advertisement is a humerous skit about a whole department of lab coated technicians who ensure that Rockwell tools “make sense”. In this ad they note that Sonicrafter’s universal fit system “fits all blades out there - unlike the other guys”.

Variable d	Value v	Z-score $P_{d,v}^{+}$
Off-Road Recreational Vehicles	True	3.396183
DIY Living	True	1.787005
Personicx Classic	Country Ways	1.748729
Home Improvement – DIYer	True	1.336364
Woodworking	True	1.258405
Hunting	True	1.180173
Personicx Classic	The Great Outdoors	1.167606
Military Memorabilia, Weaponry	True	1.107564
Personicx Classic	Full Steaming	1.080662
Personicx Classic	Acred Couples	1.068328
Science, Space	True	1.01155
Personicx Classic	Rural Retirement	1.008405
Motorcycling	True	0.956336
Auto Parts and Accessories	True	0.911406
DOB – Year	Born in the 1940s	0.908301
Motorcycle Owner	True	0.903744
Crafts, Hobbies Interest	True	0.870761

Fig. 3. Top standardized demographic scores for Rockwell Tools buyers. The score represents the number of standard deviations higher than the mean for this demographic variable compared to the US Population.

Program M_i	Corr r
20TH CENTURY BATTLEFIELDS	0.855214
FIRST WORLD WAR	0.850499
HOW THE EARTH WAS MADE	0.849239
NAPAS NORTH TO ALASKA	0.843346
FISHING W ROLAND MARTIN	0.841979
CHAMPIONS TOUR	0.836778
EUROPEAN TOUR	0.834879
BASS PRO SHOP FISHING	0.833226
HITLERS BODYGUARD	0.832817

Fig. 4. TV Programs in order of closest similarity to the customer profile.

Rotation M_i	Corr r	Rotation M_i	Corr r
MIL-M-Su + 8p-12a	0.884642	GOLF-M-F + 6p-8p	0.789733
MIL-Sa-Su + 9a-8p	0.866308	GOLF-M-F + 9a-3p	0.782984
MIL-M-F + 9a-3p	0.846683	CNBC-M-Su + 6a-9a	0.782734
MIL-M-F + 6p-8p	0.828761	DSCI-M-Su + 8p-12a	0.781056
MIL-M-F + 3p-6p	0.826187	HIST-Sa-Su + 9a-8p	0.774586
VS-Sa-Su + 9a-8p	0.818063	CNBC-M-F + 6p-8p	0.772582
GOLF-Sa-Su + 9a-8p	0.804399	MIL-M-Su + 12a-6a	0.760822
HIST-M-F + 6p-8p	0.801554	HIST-M-Su + 8p-12a	0.749103
HIST-M-F + 3p-6p	0.798157	GOLF-M-F + 3p-6p	0.746113
GOLF-M-Su + 8p-12a	0.795244	HI-M-Su + 8p-12a	0.735055

Fig. 5. Rotations in order of highest match. MIL = Military Channel, VS = Versus (a sports station), GOLF = Golf network, HIST = History Channel, DSCI = Discovery Science, HI = History International, CNBC = Broadcast Network on US television. Each rotation is a standard segment of media that can be purchased by media buyers.

Program M_i	Cost per correlated	r

	impressio n *tCPM*	
WWE FRIDAY NIGHT SMACKDOWN!	2.72	0.560
FOX AND FRIENDS SATURDAY	3.29	0.426
AMERICA'S NEWSROOM	3.37	0.253
THE FOX REPORT WITH SHEPARD SMITH	3.43	0.434
POLITICSNATION	3.91	0.422
THE RIFLEMAN	3.93	0.317
FOX AND FRIENDS SUNDAY	3.97	0.406
FOX AND FRIENDS	4.32	0.437
MAD MEN	4.45	0.555
SANJAY GUPTA, MD	4.47	0.514
CNN SUNDAY MORNING	4.48	0.491
MORNING JOE	4.53	0.348
STUDIO B WITH SHEPARD SMITH	4.75	0.232

Fig. 8. Lowest tCPM programs in the month of December 2011 where tCPM is defined as the inverse of (2). The above programs would be targeted in a campaign for the handyman product.

IV. DEMAND FINDER

The above algorithm can automatically identify a machine-generated list of top media to target. However television ad spots can cost hundreds of thousands of dollars. In order to be successful in practice, the information needs to be presented to Media Buyers in a way which empowers them to understand the available media, why it was recommended and what alternatives are available.

Demand Finder is designed to fulfill this role. It allows for interactive filtering and selection of programs based on targeting goals. The Demand Finder is a web application written in Microsoft Silverlight and utilizes its innovative PivotViewer class [9]. PivotViewer is designed to work with high cardinality discrete assets – which maps well to our domain in which we need to sort through TV programs.

The system begins fully zoomed out showing media M_i as a series of colored squares (Zoom x1; Figure 9). Green squares have high probability of buyer, and red have low probability of buyer as calculated by $r(M_i)$ (3). At a glance this shows which programs the user should explore.

The user can then zoom in to look at the high targeted programs rolling the mouse scroll wheel forward. As the user zooms in, the application begins to reveal new detail in the squares. Networks (eg. CNN, ABC) are identified using their branded logo (Zoom x2; Figure 10).

At the third level of zoom, program details become visible including targeting $r(M_i)$, cost per thousand impressions $CPI(M_i)$ (4), and impressions $I(M_i)$ (Zoom x3; Figure 11).

At the fourth level of zoom, a "heads up display" is shown which shows detailed information about the program including the demographics of the program and of the buyer target. This can help to provide information on why the program matched (Zoom x4; Figure 12). We've found that this detailed match information is extremely important for giving buyers confidence in the tool when faced with hundred thousand dollar spot buys.

Lewis's (2003) *Moneyball* presented the story of how the Oakland A's Baseball Team were able to find over-performing, under-valued players [6]. On television we aim to do the same thing. The highest value per dollar media can be revealed by selecting "sort=tCPM" and "view=histogram". tCPM is equal to the inverse of $\frac{r(P,M_i)}{CPI(M_i)}$ (2) and places each media into a value per dollar bucket (Figure 13).

An optimal plan can also be calculated by automatically selecting a set of media that maximizes (1) subject to rotation spacing rules and buyer selections. If the user selects "optimal plan = yes" filter then all airings that are in the optimal plan are shown and the remainder are filtered out. The results are shown in figure 14 (before) and 15 (after), showing the programs to purchase by network.

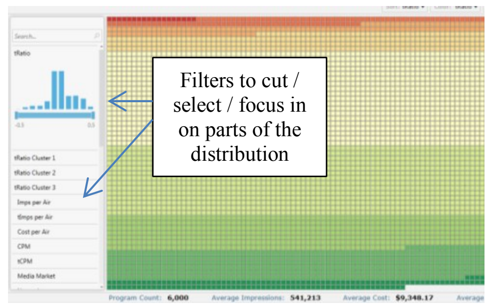

Fig. 9. Zoom x1: Fully zoomed out "Heatmap" showing programs colored by match score.

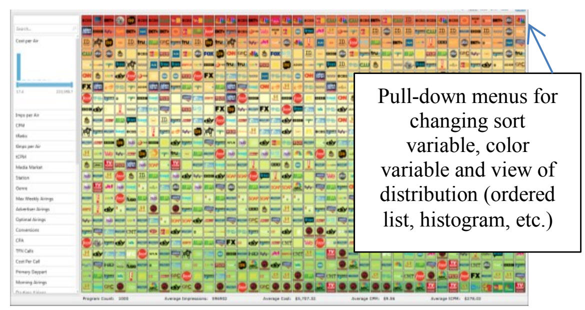

Fig. 10. Zoom x2: Zooming in reveals station icons, but still colored by match score.

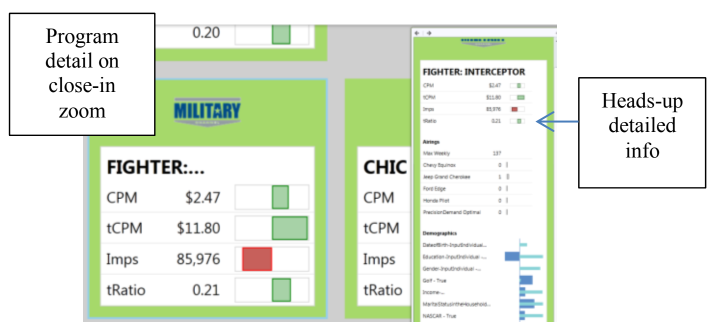

Fig. 11. Zoom x3: Zooming further brings up details of the programs including CPM, tCPM, Imps, tratio (targeting score). Handyman buyers tend to watch "Fighter Interceptor" on Military Channel.

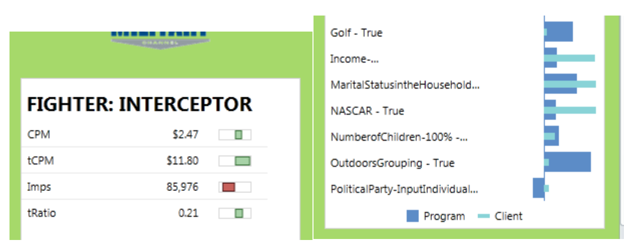

Fig. 12. Zoom x4: The Demographics of the program and advertiser target are shown in the Heads Up Display overlay. Wide bars refer to the standardized demographic score for the program. Narrow bars are the standardized demographic score for the advertiser's target. The program

being investigated is a good match on Marital Status, NASCAR interest, Outdoor interest and so on.

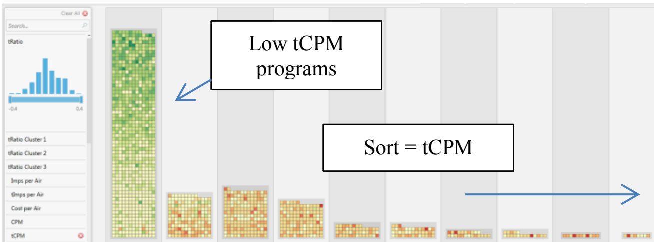

Fig. 13. The best value per dollar programs can be shown by selecting “Sort by tcpm” and “histogram”. The best value per dollar programs are towards the left histogram bin.

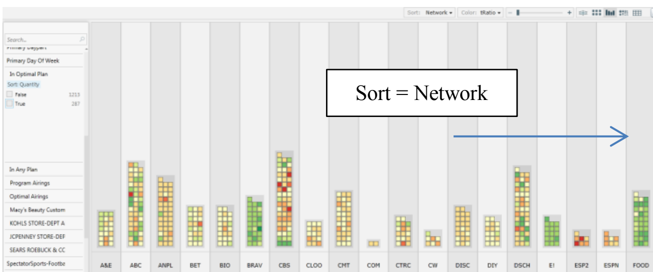

Fig. 14. TV programs by network.

Fig. 15. Switching on a filter for “optimal plan = yes” shows only the TV programs that should be purchased given the user’s budget. The above view shows optimal programs by network. This makes it easy for buyers to contact the stations

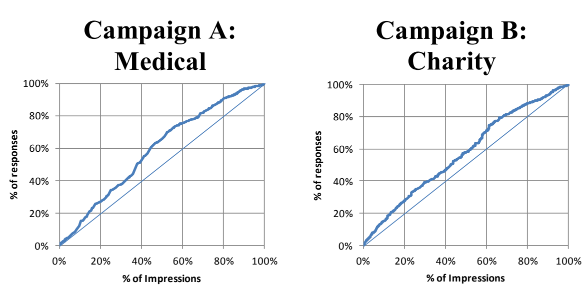

Fig. 16. TV Campaign Results from Targeting System. Graphs show cumulative phone response distribution (y-axis) versus targeting score (x-axis) for two real television campaigns. Approximately 50% lift is generated in both campaigns from the top 20% of targeted programs.

V. Results

The Set Top Box Television Ad Targeting algorithms that we have described in this paper have been tested in practice in several live television campaigns [3]. Figure 16 shows results from two campaigns which embedded phone numbers in the advertisement. Campaign A (“Medical”) was an advertisement for a discretionary medical procedure. The campaign comprised of over 11,761 airings with a spend of 2.46 million dollars. Campaign B (“Charity”) was an advertisement for a charity in Africa to provide medical care for millions of people. This campaign consisted of 1,663 ad airings with a spend of 145,363 dollars.

In both campaigns we sorted airings by the targeting score $r(M_i)$ descending, and plotted cumulative phone responses. In the top 20% of airings, approximately 30% of responses were isolated. As a result, in both cases we observed about 1.5x phone response lift for selecting the top 20% of targeted airings [3].

VI. Conclusion

Demand Finder has proven to be invaluable for (i) making the concept of TV targeting easy to understand for lay users (ii) building optimal TV media plans, (iii) understanding why programs match or may not match, and (iv) showing live campaign performance. The visualization is built on the idea of being able to filter, sort, and select discrete, buyable, media (often programs) which are rendered as squares that can be manipulated by the user. This results in an intuitive and compelling visual experience. This allows Media Buyers to quickly navigate the TV landscape and identify the highest targeted, least cost media targets for their campaign – a process which in previous live campaigns has been able to deliver 50% higher revenue lift.

Acknowledgments

Thanks to Dyng Au, Brian Burdick, Jon Borchardt, Amanda Powter, Todd Otis, Huseyin Ulger, John Crim, John Sobieski and Dana Chapman.

VII. References

[1] C. Dawson, “Television Advertising: In Need of Reinvention?”, International Journal of Marketing, Vol. 16, No. 4. 1996.
[2] C. Kennedy, M. Loughney and A. Waterston, “ABC TV – Untethered TV”, Re:Think 2012, The Advertising Research Foundation: New York. 2012.
[3] B. Kitts, D. Au, B. Burdick, “A High-Dimensional Set Top Box Ad Targeting Algorithm including Experimental Comparisons to Traditional TV Algorithms”, Proceedings of the Thirteenth IEEE International Conference on Data Mining, Dallas, TX, IEEE Press. 2013.
[4] B. Kitts, L. Wei, D. Au, S. Zlomek, R. Brooks, B. Burdick, “Targeting Television Audiences using Demographic Similarity”, Applications of Data Mining and Modeling in Government and Industry Workshop (ADMMGI2010) in Workshop Proceedings of the Tenth IEEE International Conference on Data Mining, IEEE Press. 2010.
[5] B. Kitts, B. Burdick, D. Au, T. Roberts, “Method and System for Automatically Targeting Ads to Television Media Using Demographic Similarity”, USPTO Patent Application No. 13/209,346, Filed 2010.
[6] M. Lewis, Moneyball: The Art of Winning an Unfair Game, W. W. Norton and Company. 2003.
[7] A. Segal, “Nielsen Ratings: An Inaccurate Truth Out of date television ratings system exposed”, http://cornellsun.com/node/23180, The Cornell Daily Sun, April 26, 2007
[8] IAB, Platform Status Report: Interactive Television Advertising, 2012.

[9] Microsoft Silverlight PivotViewer, http://www.microsoft.com/silverlight/pivotviewer/ accessed Oct 2013.