WP7:
Ad Recommendation Engine

Realeyes
Elnar Hajiyev, CTO
Gabor Szirtes, Head Scientist
Pierre Lupi Chen, Data Scientist

sowa
Automatic Sentiment Analysis in the Wild

Realeyes

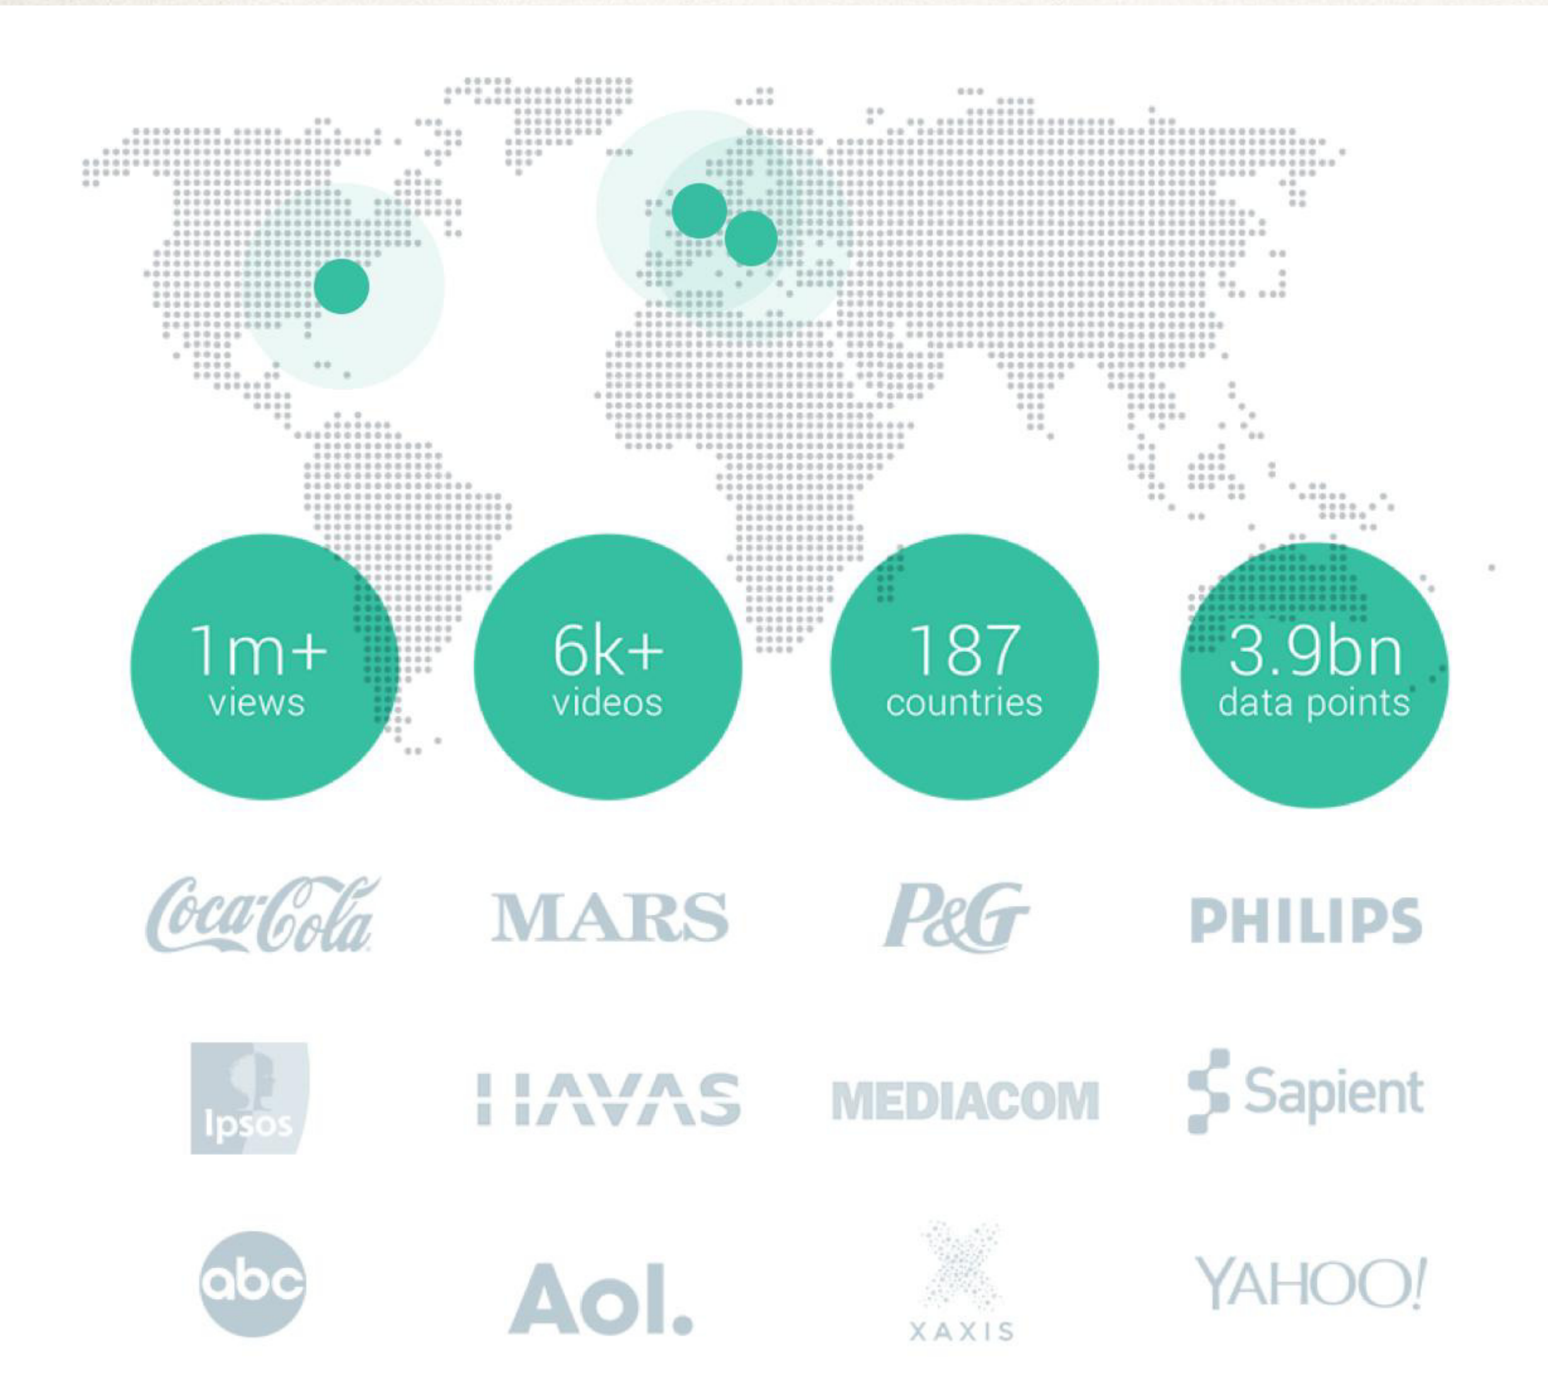

Founded at Oxford University

40+ top professionals in London, Boston and Budapest offices, 10 PhDs in R&D

Global technology leader in emotion measurement via webcam

Analytics value from largest emotions database linked to real-life outcomes

Growing SaaS sales to Brands, Agencies, and Media Companies

sewa
Automatic Sentiment Analysis in the Wild

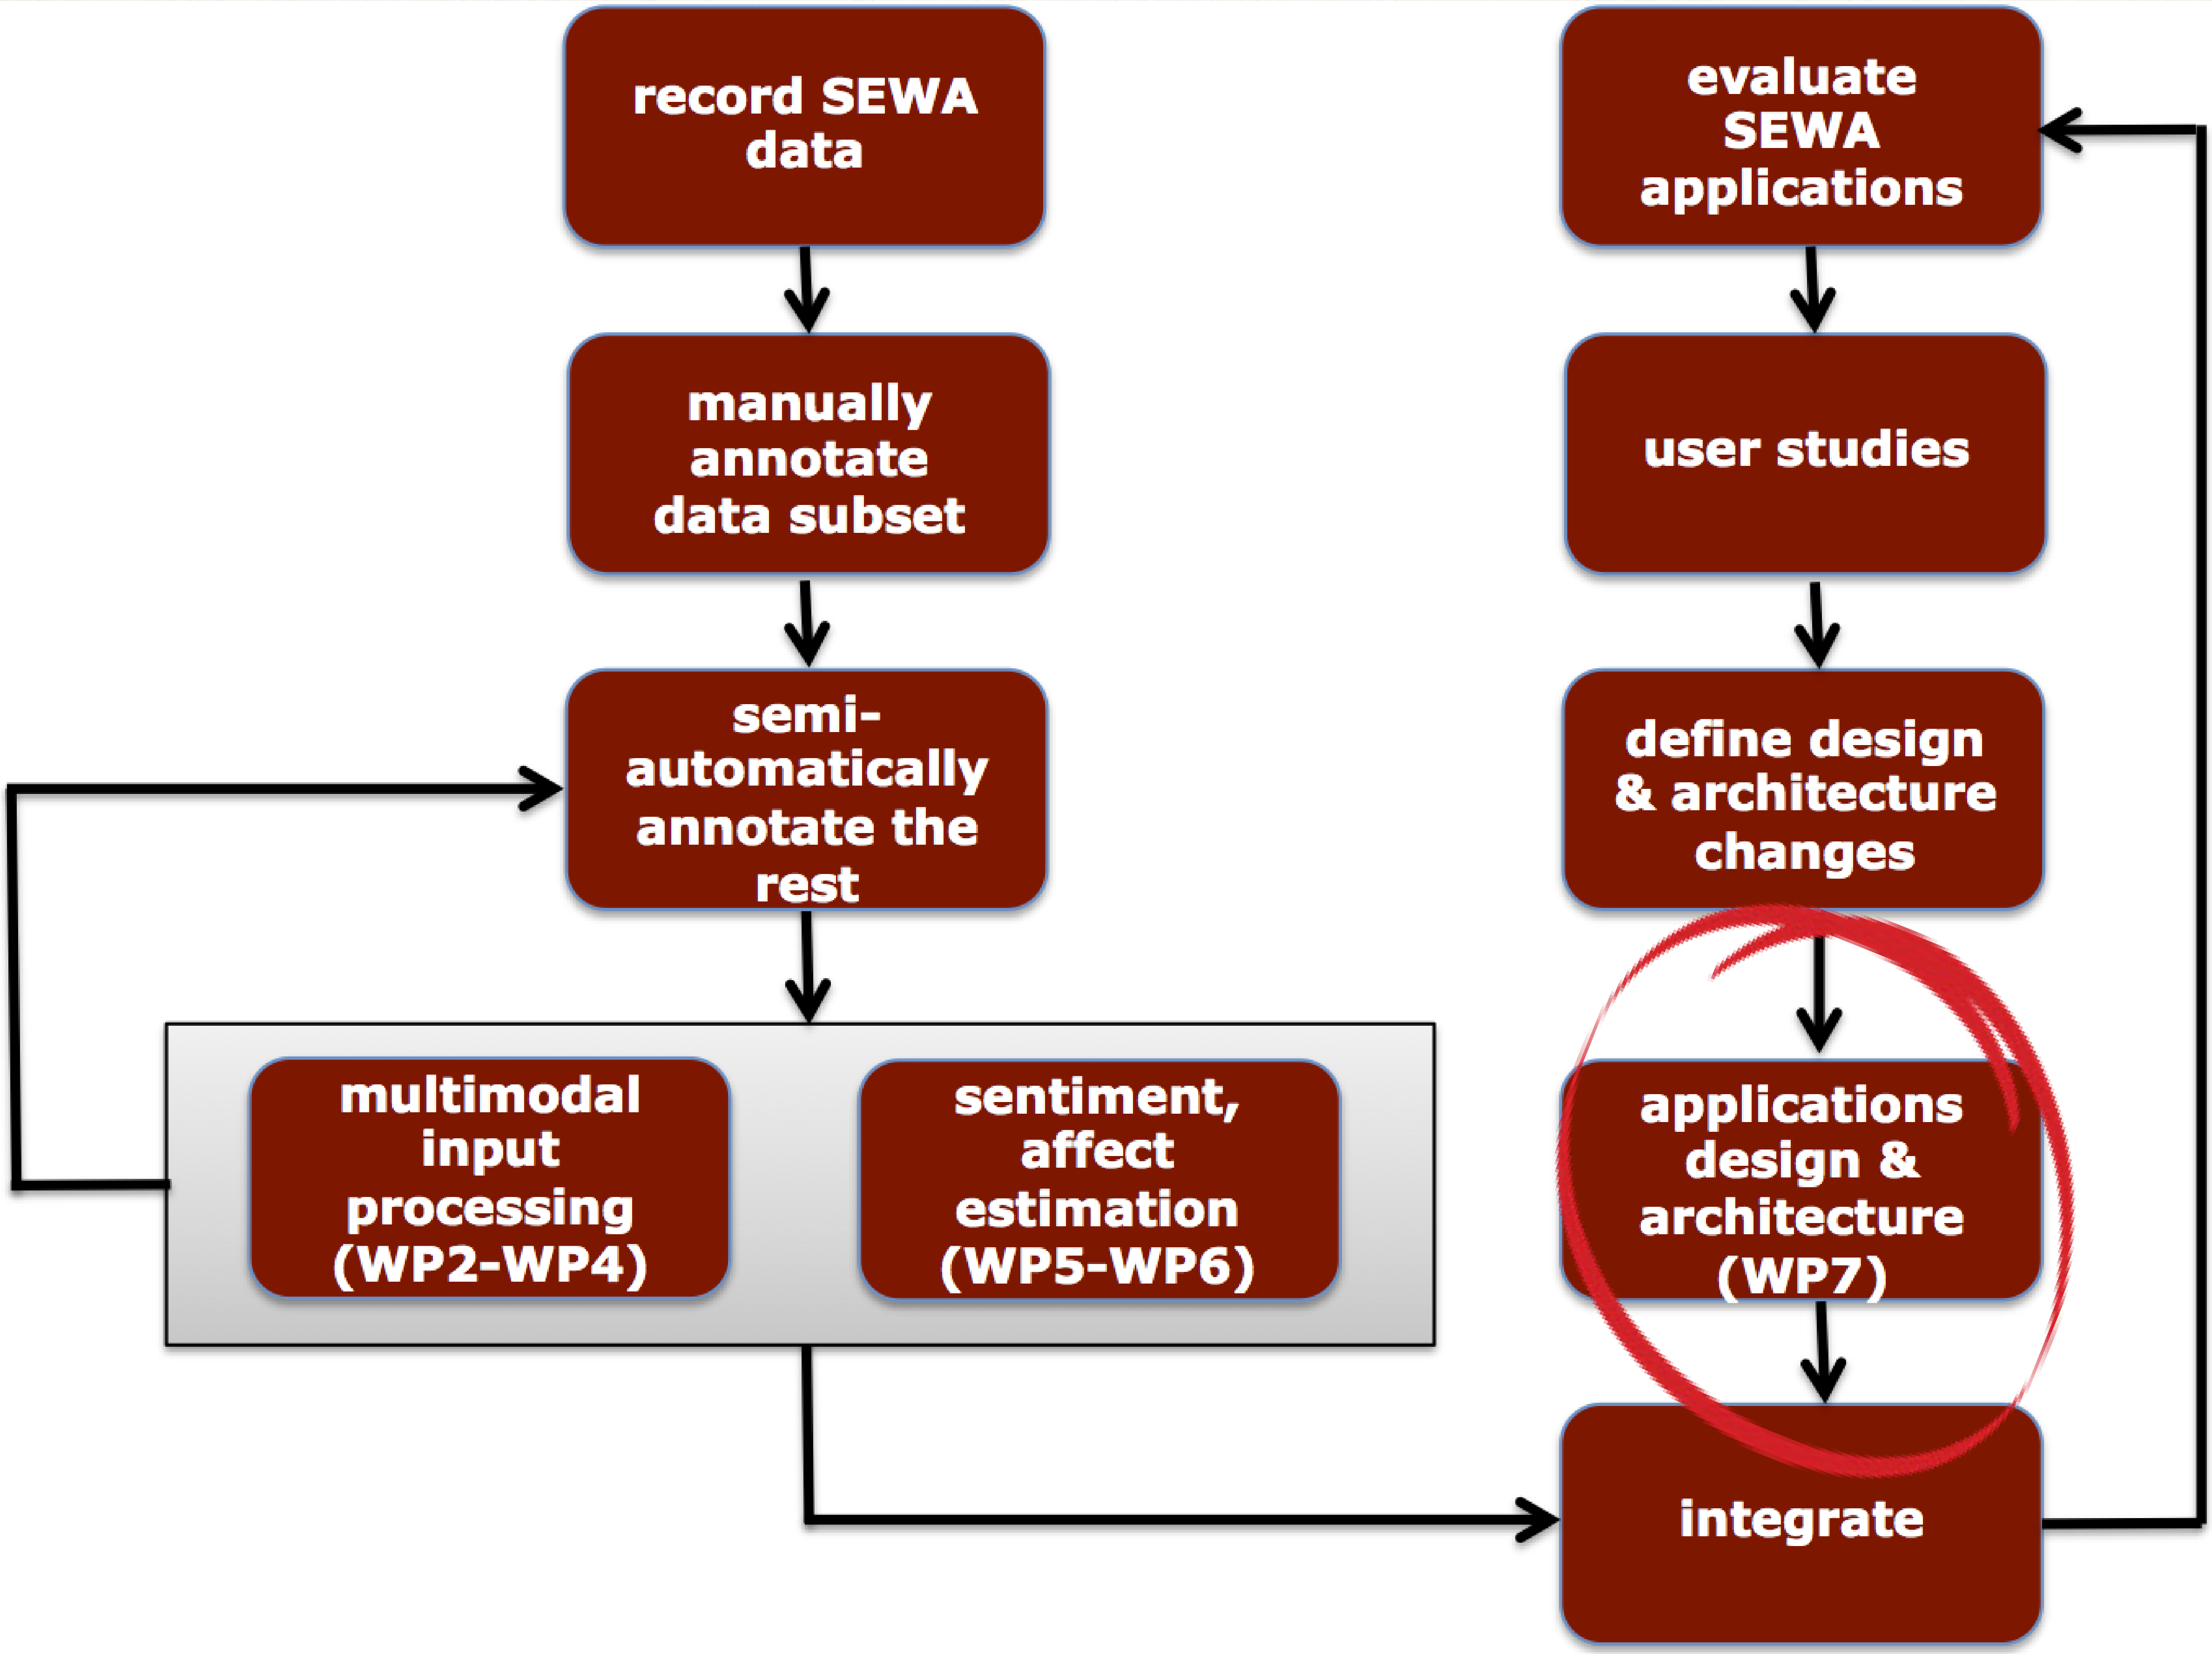

sewa

Automatic Sentiment Analysis in the Wild

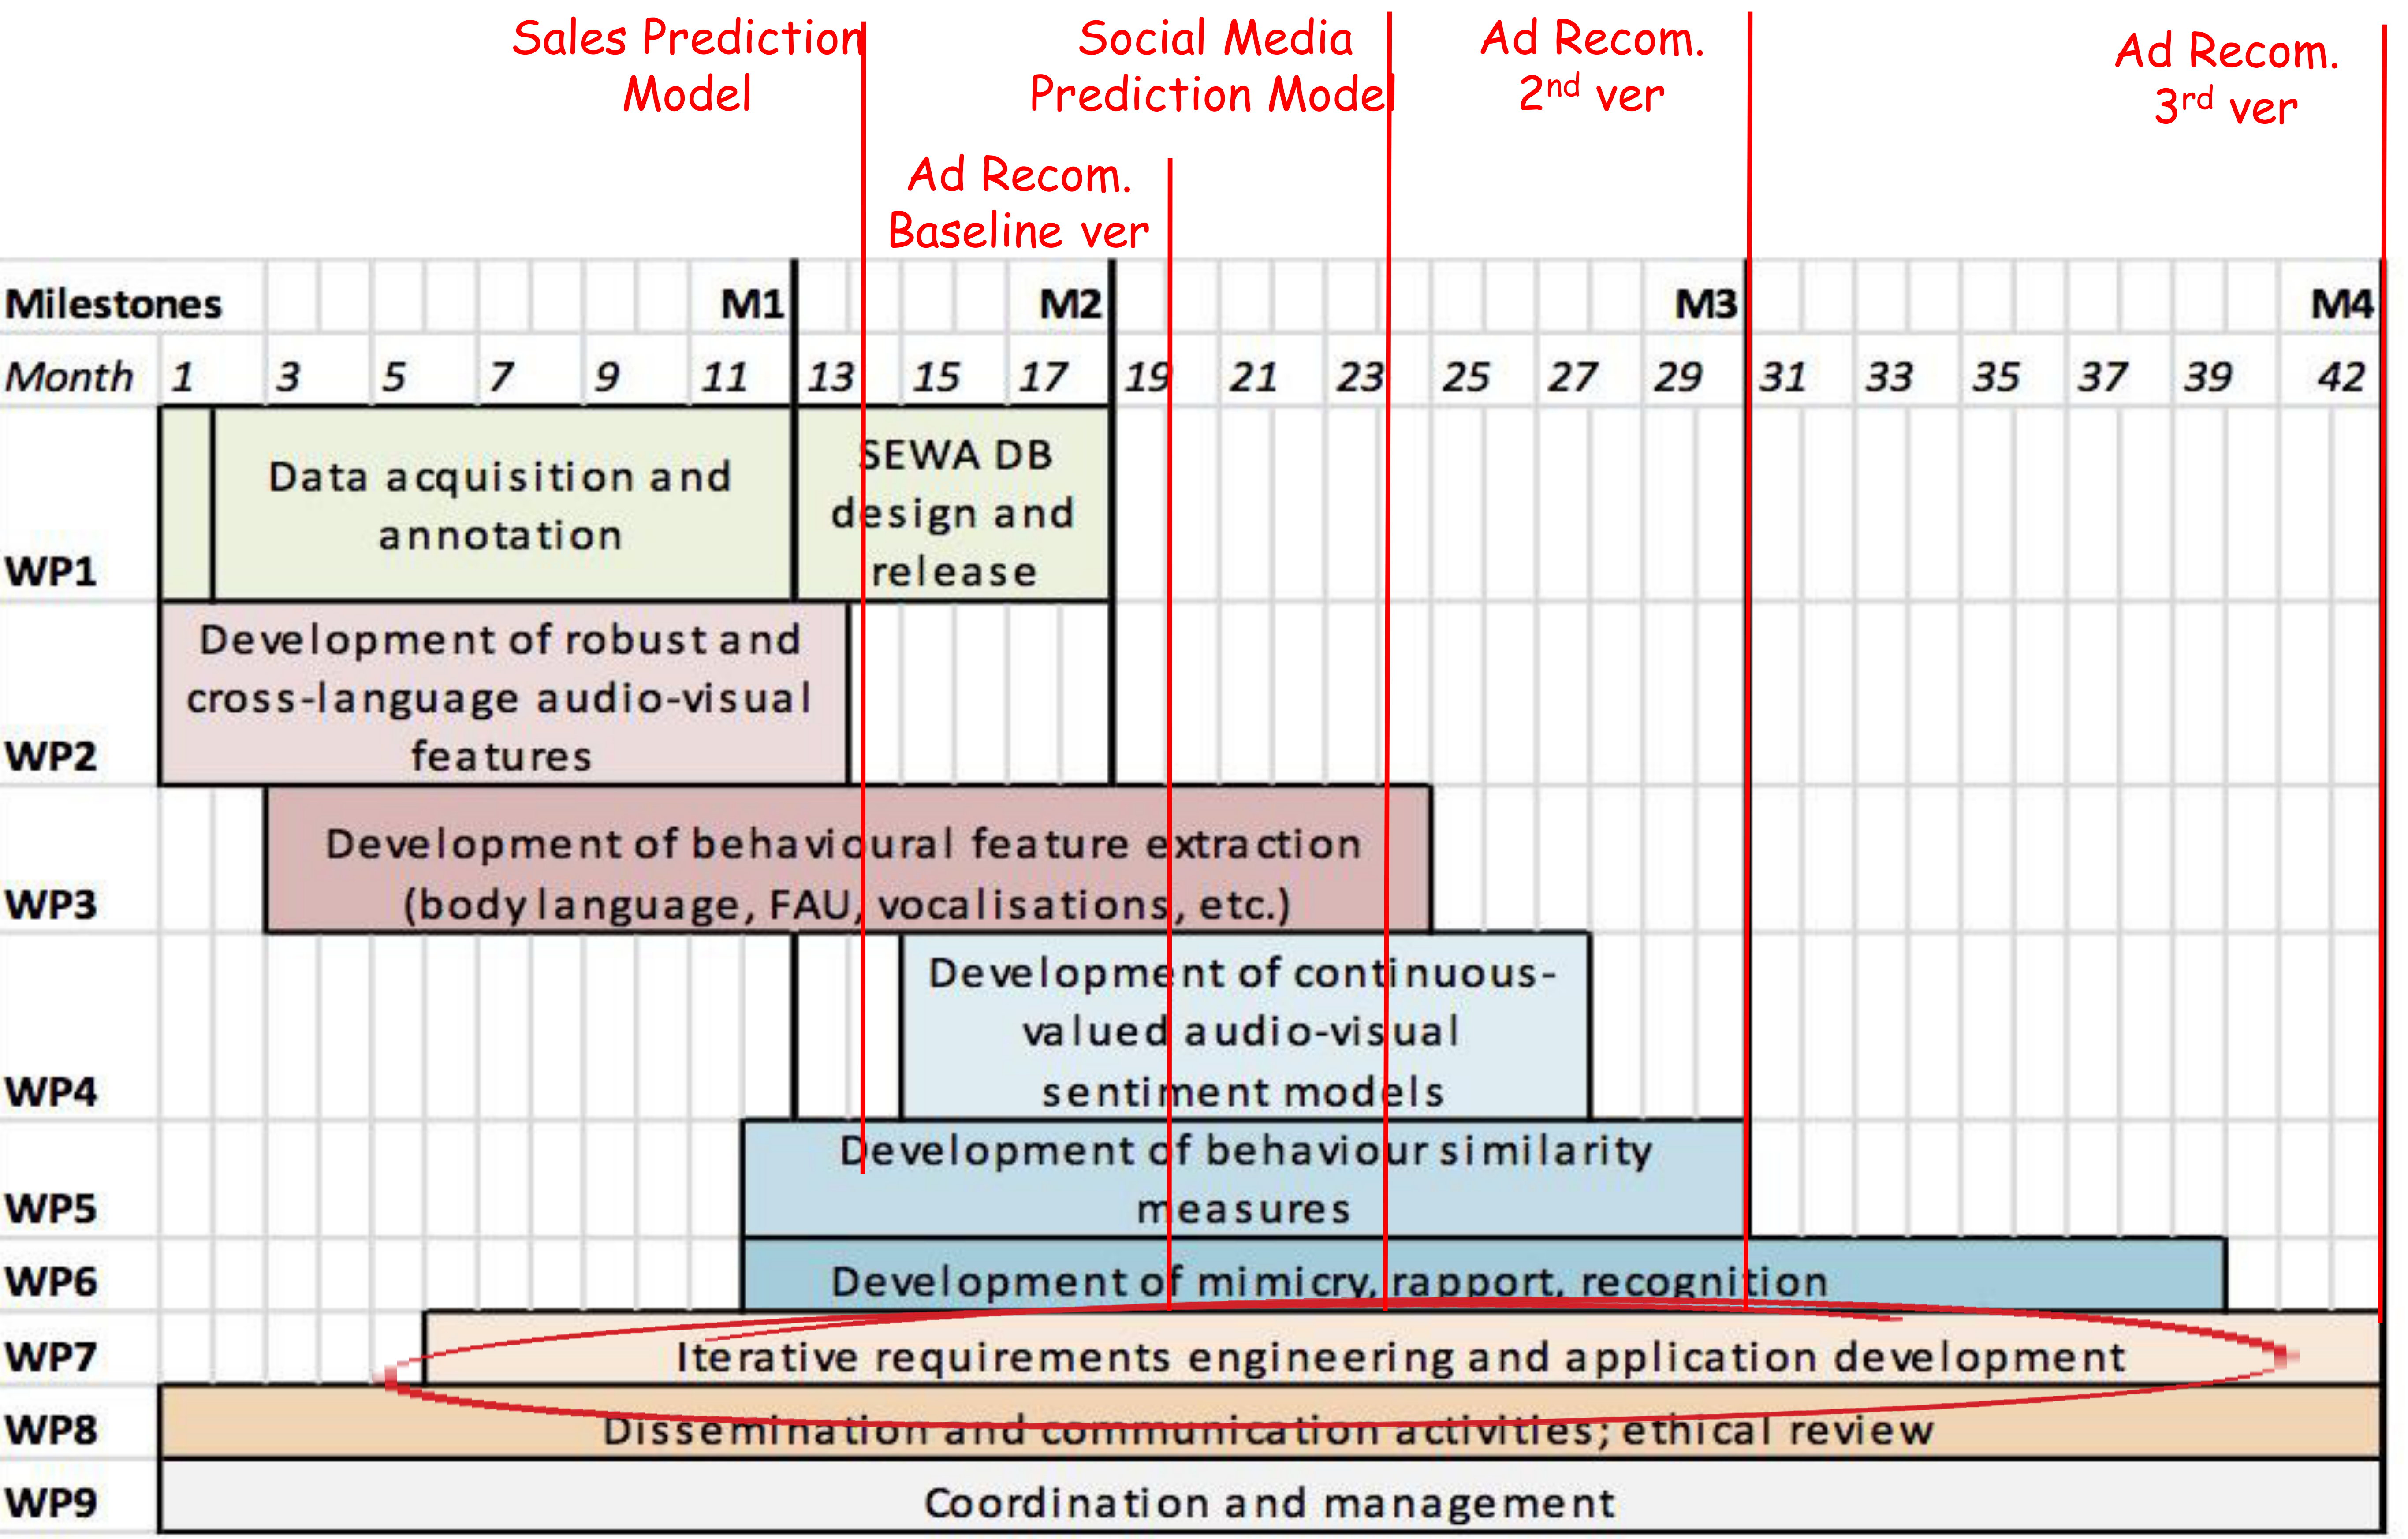

Objectives of WP7

- Report on user requirements for each SEWA application (D7.1)
- Initial version of the Ad Recommendation Engine (D7.2)
- Second version of the Ad Recommendation Engine (D7.4)
- Final version of the sentiment-driven Ad Recommendation Engine (D7.6)

User requirements

- Purpose of the ad recommendation engine:
  - Enable use emotional and behavioral information to show right ads to the right audience
- Why digital video advertising?
  - Projected spend $28.08 billion in 2020 in US alone*
  - Core area of expertise for Realeyes, 5+ years experience
- Target user groups
  - The advertiser (brand owner)
  - The consumer (audience)
  - The publisher (content owner)

* http://www.emarketer.com/Article/Digital-Video-Advertising-Continues-Expand/1013722

User requirements

- How is it contributing to the existing industry methods?
  - Can be used for pre-testing to drive better targeting and ad design improvement
  - Richer second by second data, allowing better impact understanding
  - Links ad impressions with user impact and gives information about attitudinal impact of the ad
  - Fills the measurement gap with brand awareness campaigns
  - Fast and inexpensive (compared to similar methods, e.g. EEG)
  - Matching emotional level of the content with ads

User requirements

- Solution development requires partner participation:
  - Brands:
    - Ultimate benefactor
    - Own crucial sales or social media data
    - Fully aware of any marketing campaigns or seasonality effects
    - Can help measure recommendation effectiveness
  - Data Management Platforms (DMPs):
    - Know what marketing campaigns are being executed
    - Fully aware and driven by the challenges of the target user groups
    - Compete to get higher quality data
    - Can help measure recommendation effectiveness

Ad Recommendation Engine

- A general recommender engine **predicts** a **score** that a **user** would add to an **item**. Items with the highest scores are then recommended to the user.
- Prediction can be based on matching **similarities** between users and items and on known preferences of the users (history): **integrated collaborative filtering approach**
- Our goal: recommend –show– maximally **relevant ads** to **viewers**.

Ad Recommendation Engine

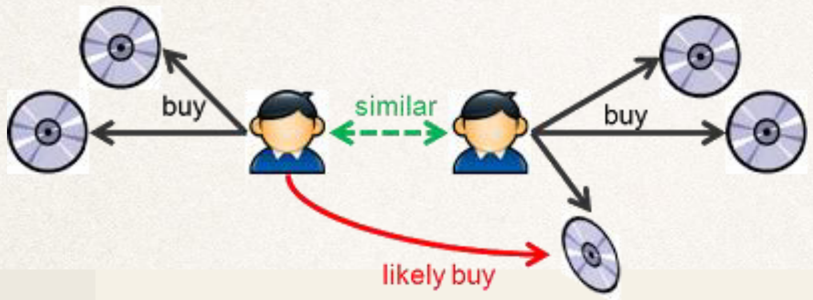

User similarity based

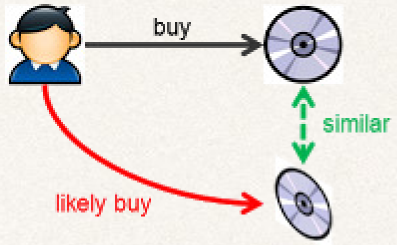

Item similarity based

[http://horicky.blogspot.hu/2011/09/recommendation-engine.html]

Challenges

- No individual rating or action available as a score for relevance
- Difficult to define similarities between ads
- Difficult to connect viewers’ attributes with ads

Our solution

- **Relevance** can be defined by sales lift scores (ad's contribution to sales)
- Interaction between ad and viewer can be measured via **emotion responses**
- **Recommendations** can be made **at group level** (for user segments)
- **Viewers** can be clustered by similarities using **past emotion responses** and external attributes (tracking data, **meta-info**)
- **Ads** can be clustered by similarities of **emotion response profiles** (aggregate group responses)

sewa

Automatic Sentiment Analysis in the Wild

Our solution – Attributes

Emotion response at panel level

Meta-info

Info on **panelists**:
Past emotion response
Tracking data
Meta-info

Ad

Attributes:

- Emotion responses
- Tracking data (history, demographics, location)
- Meta-info (personality, interest)

Viewer

Tracking data

Meta-info

Past emotion response

real eyes

Our solution – Audience Selection

Step #1 Train a model that can predict sales performance of an ad (score) from emotion responses of viewers.

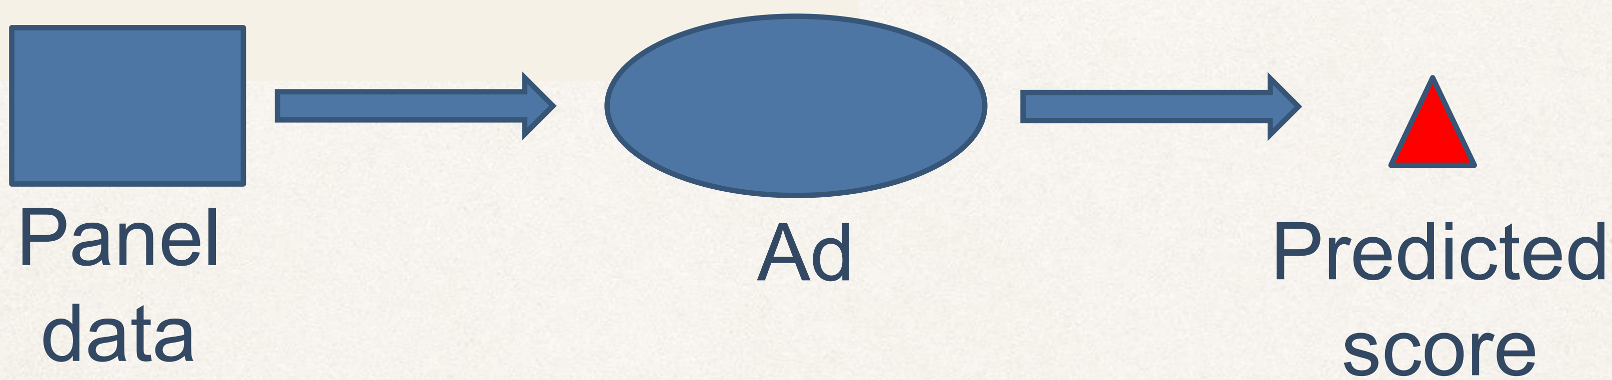

Such model can be used for recommendation of whether to air or not air the ad based on panel response data.

Our solution – Audience Selection

Step #2 For a given ad identify those viewers whose emotion responses would yield the highest predicted score.

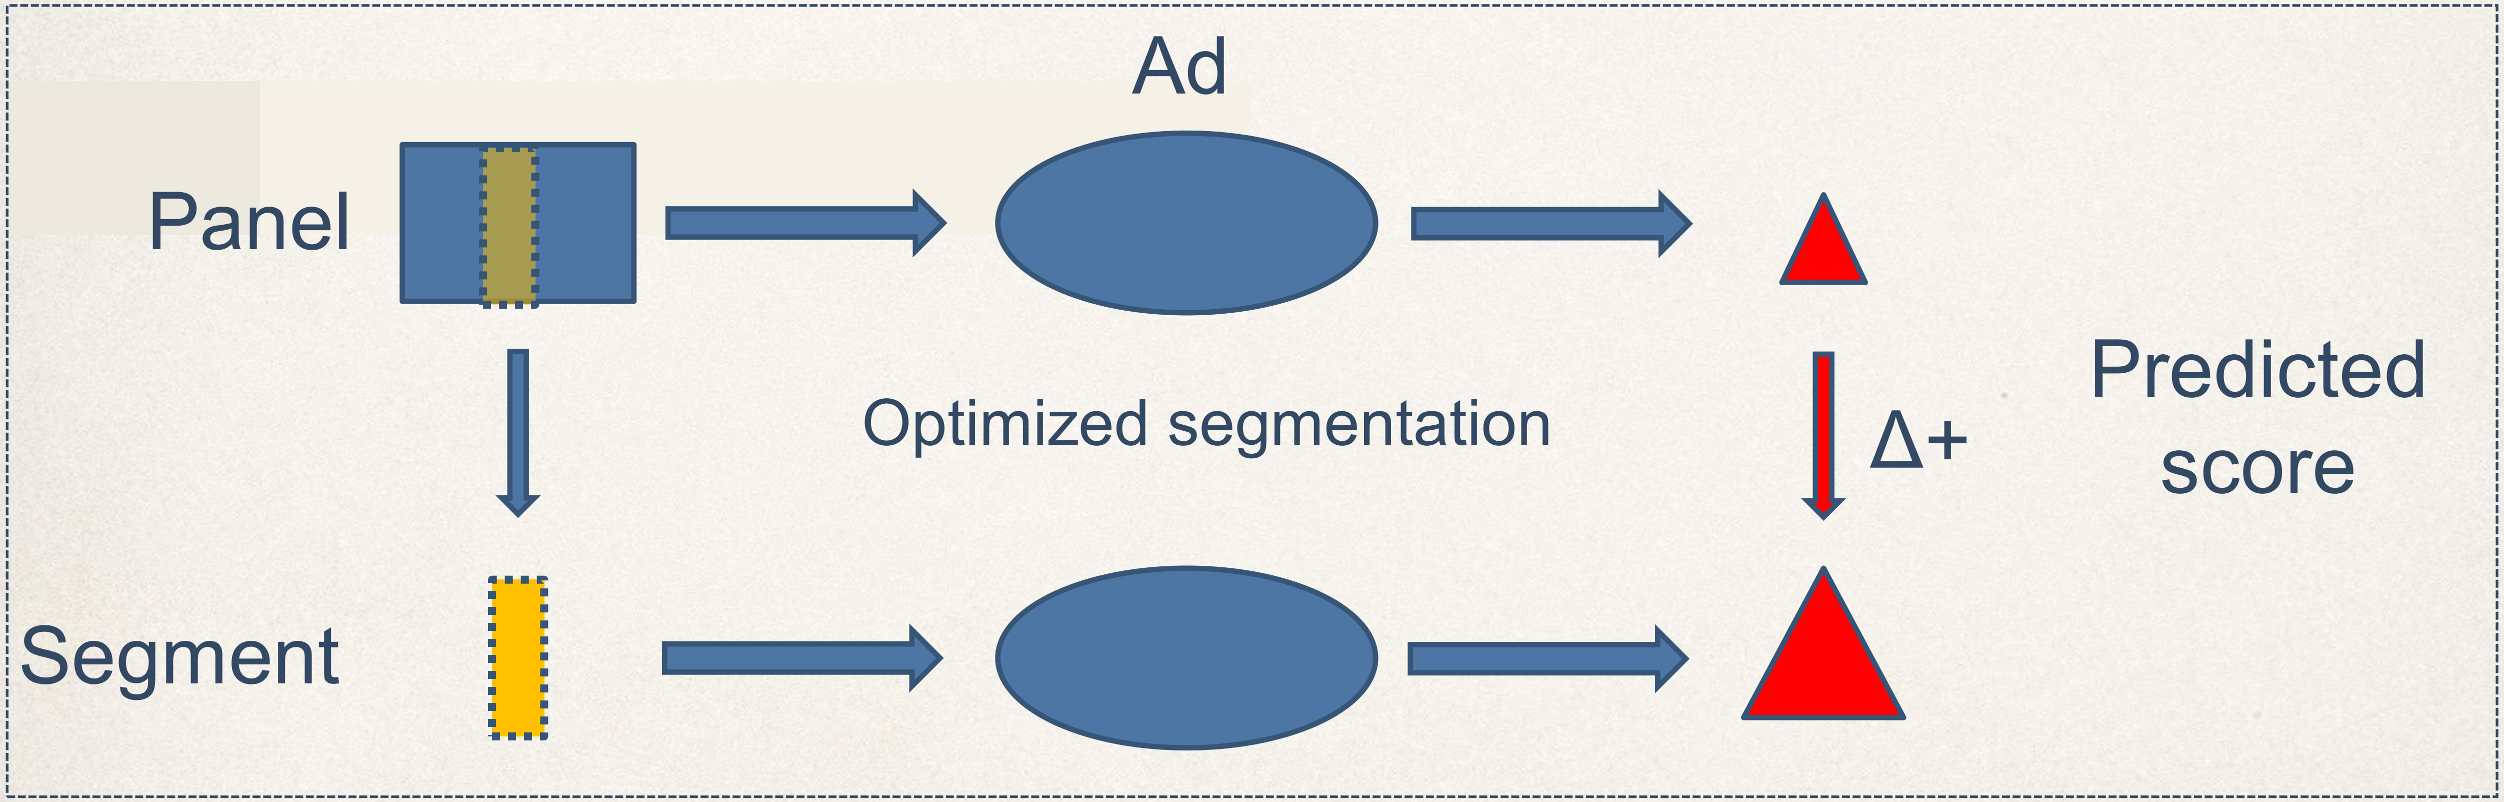

Our solution – Audience Selection

Step #3 Recommend the ad to viewers most similar to the ones selected above.

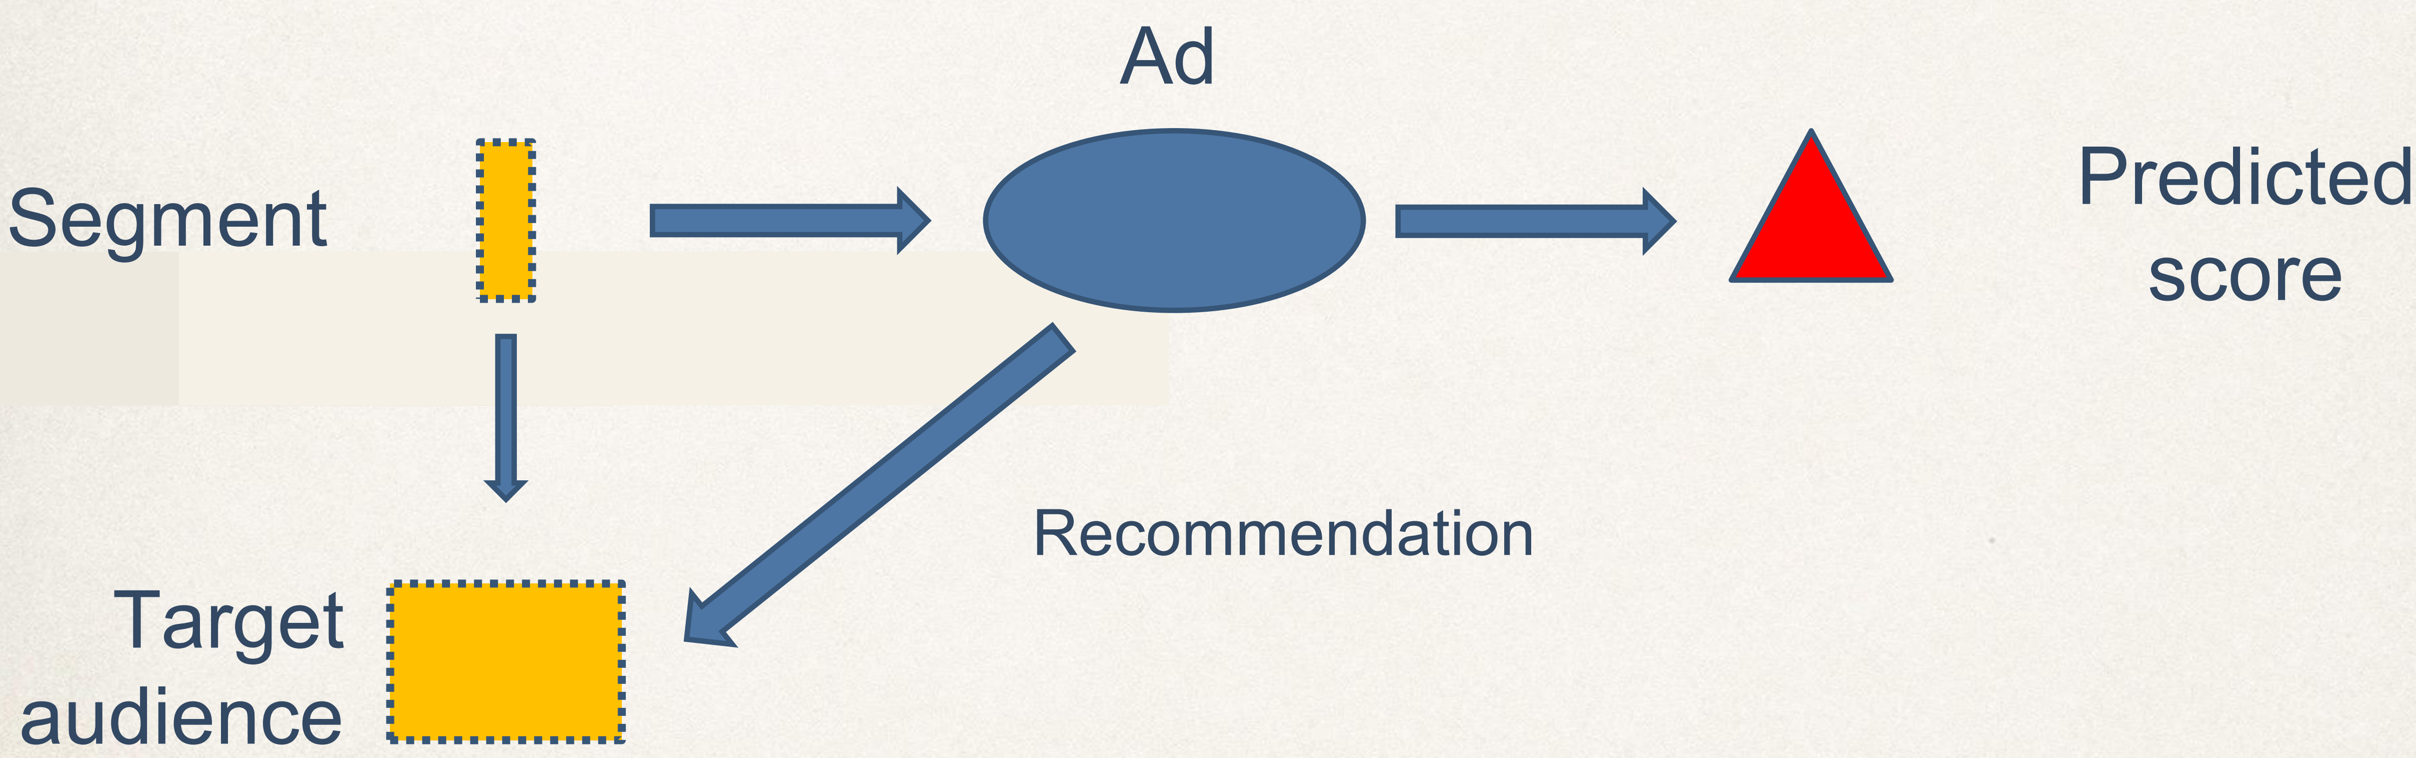

Our solution – Audience Selection

Step #4 For viewers with past measured emotion responses recommend new ads

that will likely elicit emotion responses yielding high predicted scores.

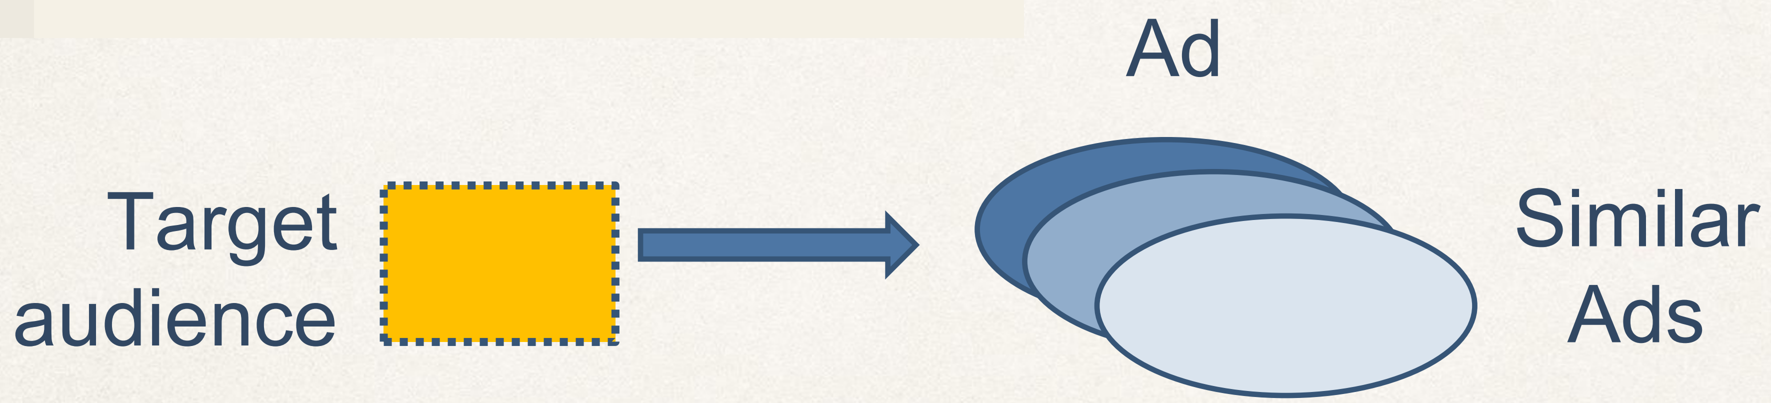

Summary

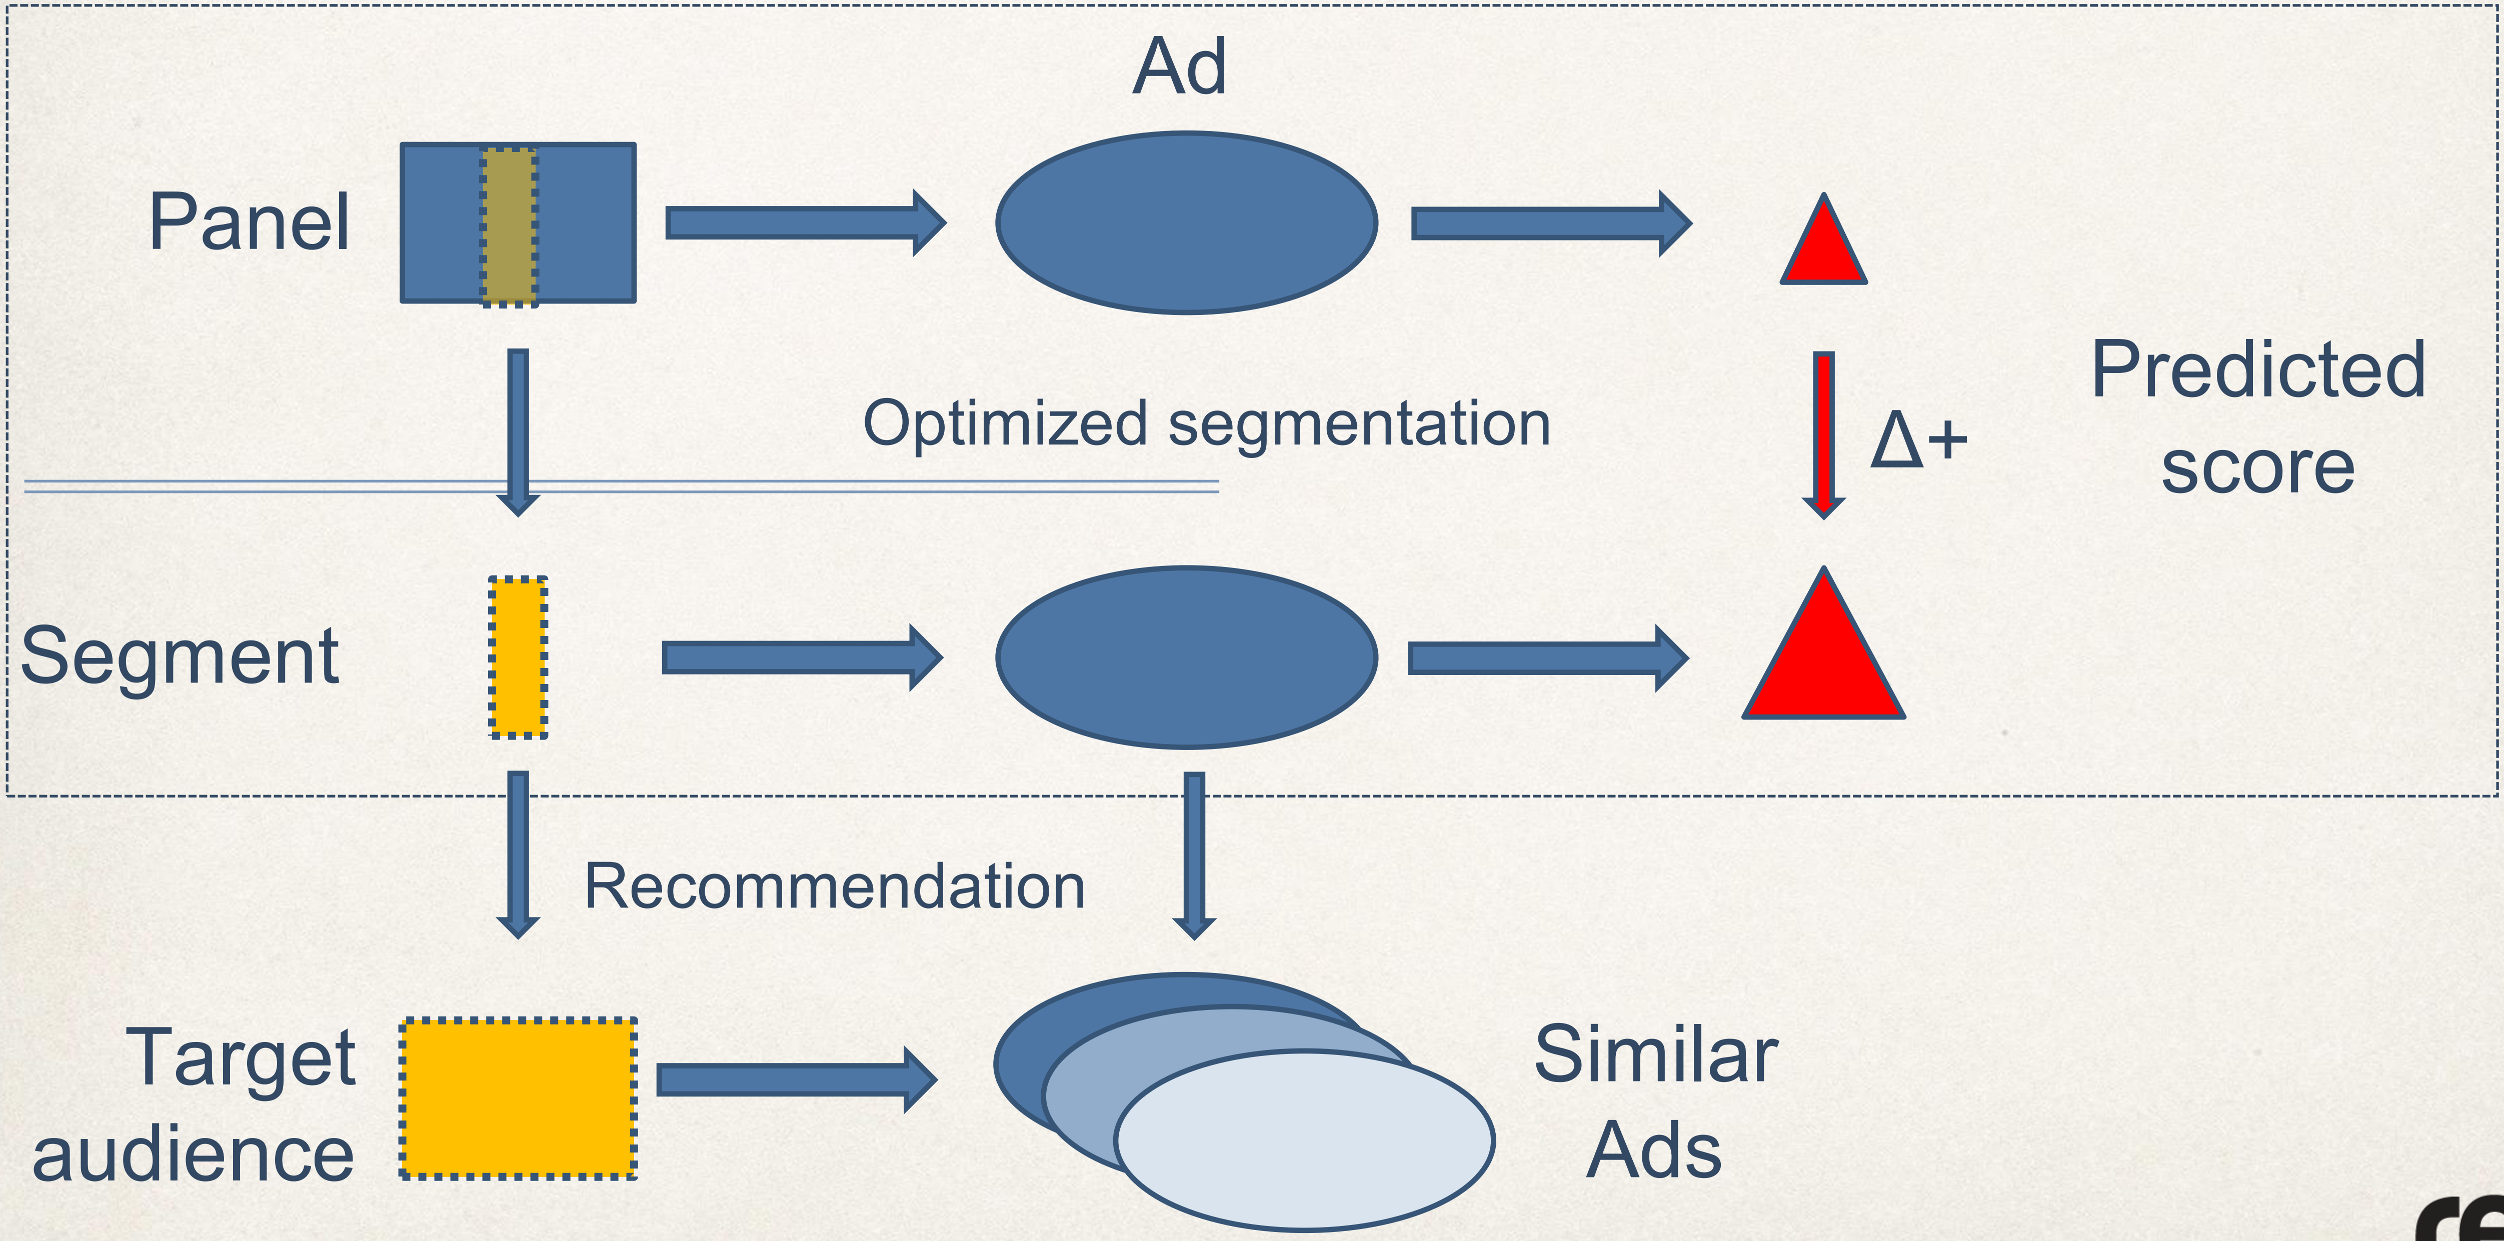
Ad
Panel
Optimized segmentation
Δ+
Predicted score
Segment
Recommendation
Target audience
Similar Ads

Step #1: build a sales lift prediction model (M2-M14)

Task #1: Collect emotion responses from panel to ads with available sales lift data (provided by MARS)

Task #2: Exploratory Data Analysis for optimal data representation

Task #3: Create, fine-tune and validate predictive models.

Constraints: cost, speed, transparency

Results #1: Data Collection

Data collection:

- We have built an online, scalable, cloud-based, dynamic data collection platform that allows for cost optimization, demographics tracking and multi-modal data collection.
- World's largest emotion data (discrete facial expressions) linked to sales information (12k viewers, 149 ads, 42k views, 6 countries, 4 categories)

Results #2: Signals

Exploratory Data Analysis:

- Created hand crafted representations based on existing shape alignment methods and facial expression classifiers (benchmark for SEWA integrated versions)

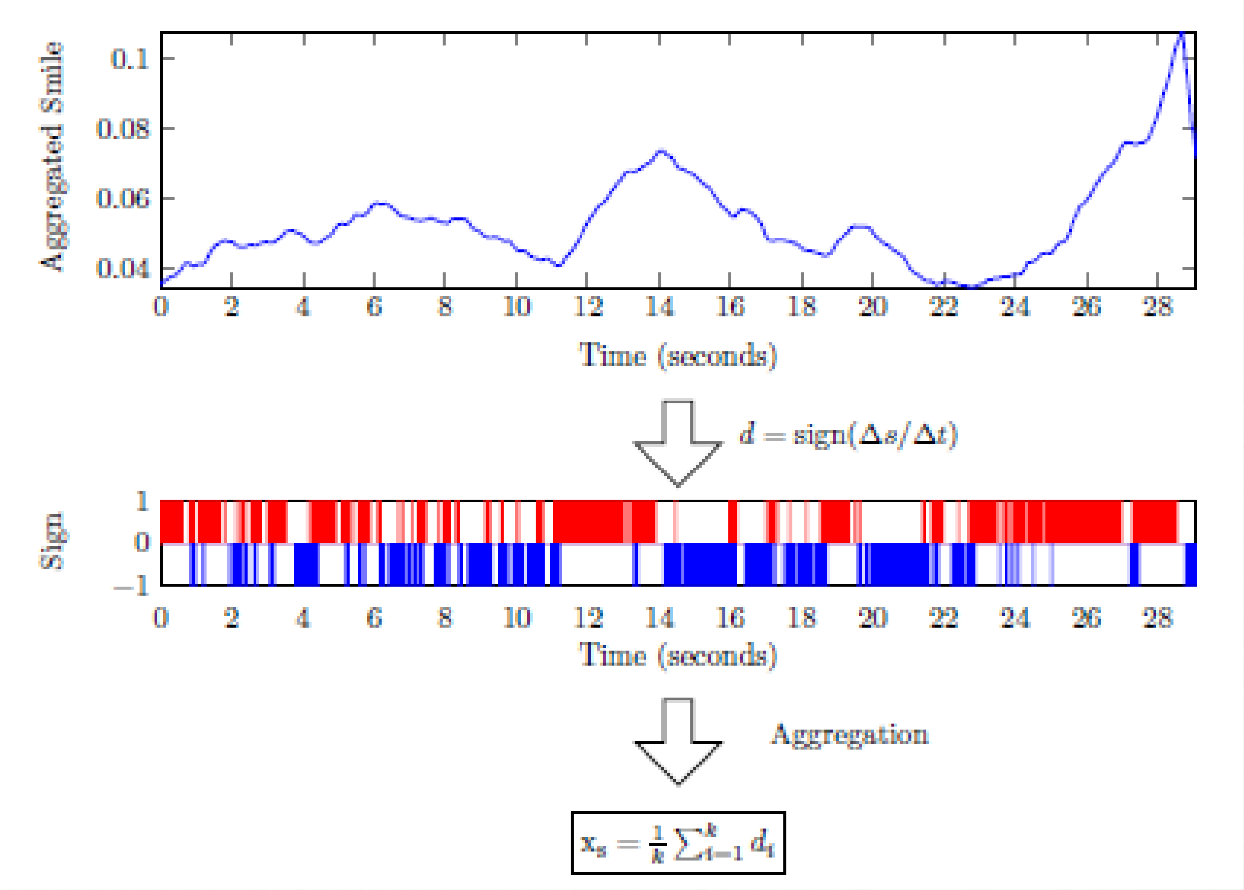

"Smile Dynamics" signal

Results #3: Modelling

- We have built a low complexity ensemble model of simple linear regressors that can classify ads as high vs low performing in sales terms. Non-linearity is introduced via thresholding
- Modeling and results are submitted to the Special Issue of Image and Video Computing and presented at I-COM

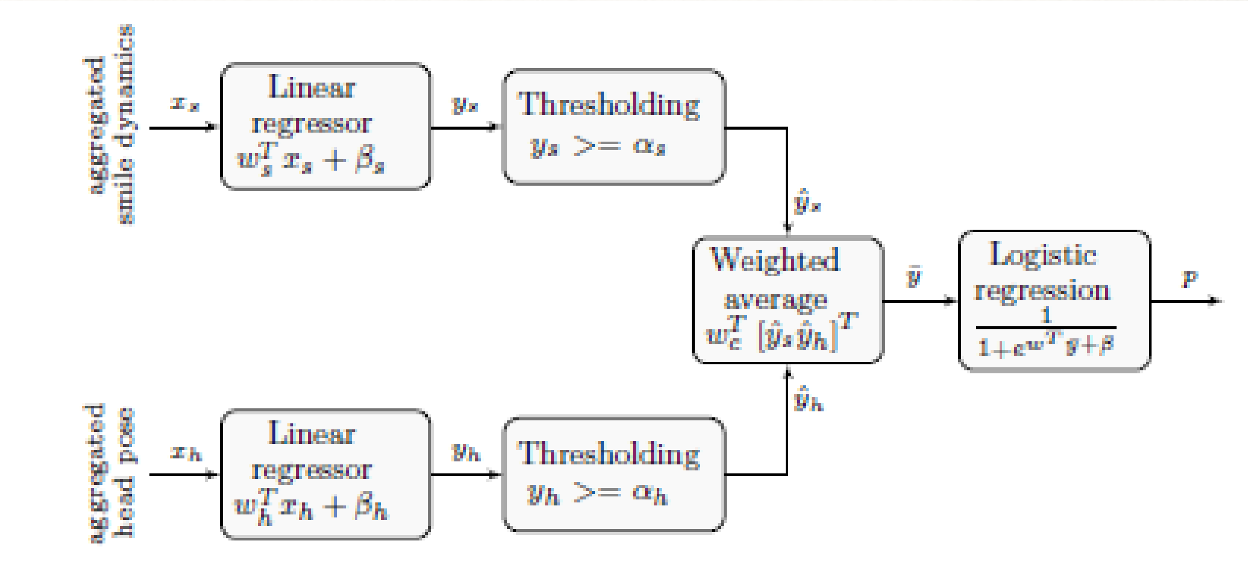

Results #4: Validation

We have validated the obtained model by traditional k-fold Cross-Validation as well as by “Leave One Label Out” type validation scheme (to check robustness against factors not considered in the model like region or product category)

repeated 10-fold CV	Accuracy	ROC AUC
Dynamic model	72.7±2.3%	0.74±0.03
Random model	52.3±2.8%	0.50

Category	Acc.	ROC AUC	#low	#high
Chocolate	69.6%	70.0%	24	22
Food	75.0%	70.0%	10	2
Petcare	72.4%	69.4%	37	21
Wrigley	75.8%	70.2%	13	20
average	73.2%	69.9%		

Region	Acc.	ROC AUC	#low	#high
Australia	59.3%	68.8%	18	9
France	86.7%	85.7%	8	7
Germany	63.6%	67.7%	10	12
Russia	81.8%	71.4%	15	7
UK	79.4%	75.9%	20	14
USA	79.3%	82.0%	13	16
average	75.0%	75.1%		

realeyes

Results #5: Product extension

Select media

Norms | Chart views | Generate video | Export | Share

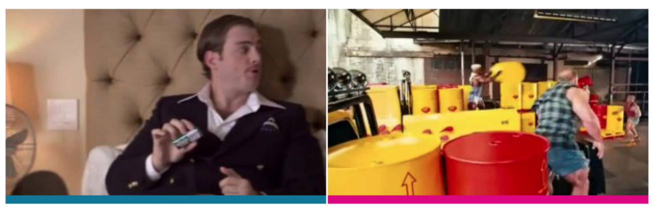

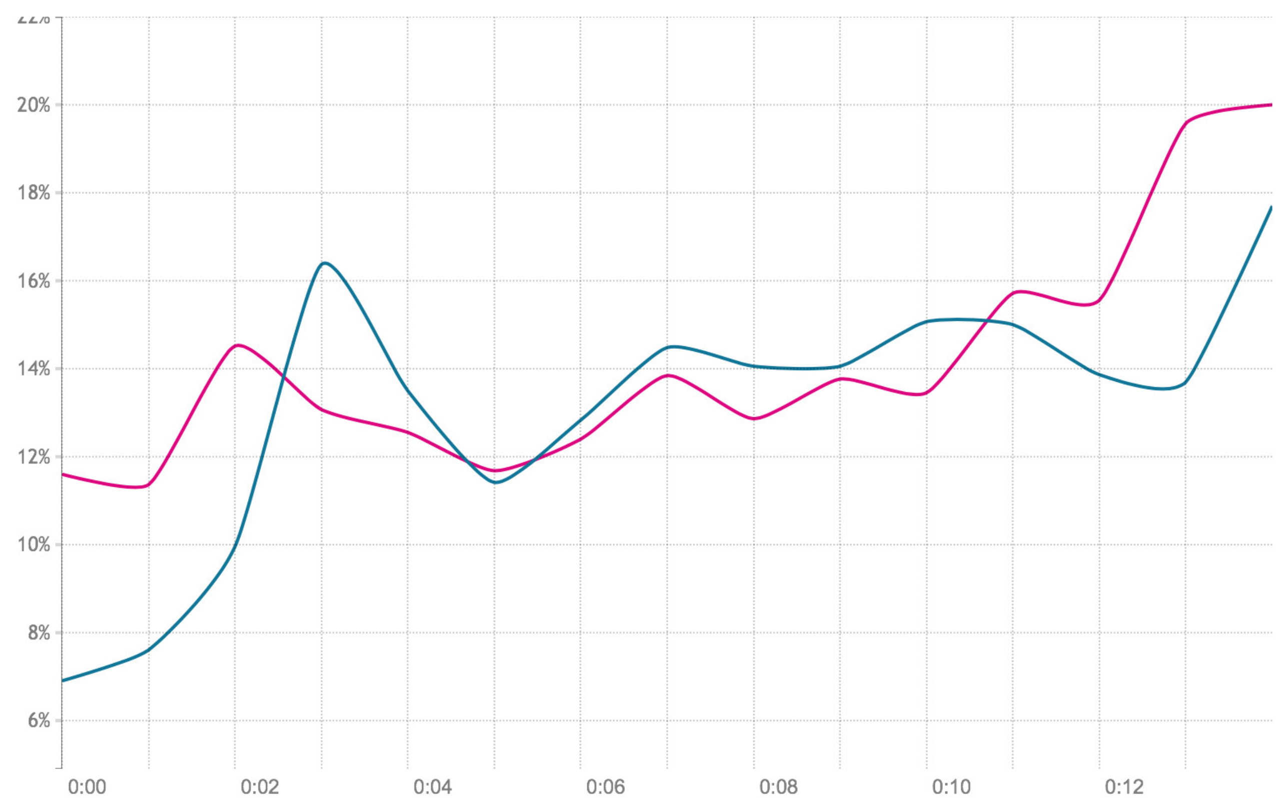

Emotions

Happy | Confused | Disgusted | Sad | Scared | Surprise | Engagement | Negative | Valence

Metrics

% of people | Average | Max | EmotionAll ® | Predicted Sales

Media

Sum | Compare | US Wrigley 2... | AU Wrigley 2...

Gender

Sum | Compare | Female | Male

Age

Sum | Compare | 29 or less | 30-49 | 50 or over

Country

Sum | Compare | Australia | United State...

MEL build master-101 master-101	Views	Prediction	Probability	Happy
AU Wrigley 2013 Eclipse Chewy Mints Chew your way to fresh 15	263	Low Sales Lift	34%	13.1% 14.5%
US Wrigley 2014 Starburst Land of Intensity 15	299	High Sales Lift	53%	14.1% 11.4%

realeyes

sewa
Automatic Sentiment Analysis in the Wild

Results #5: Product extension

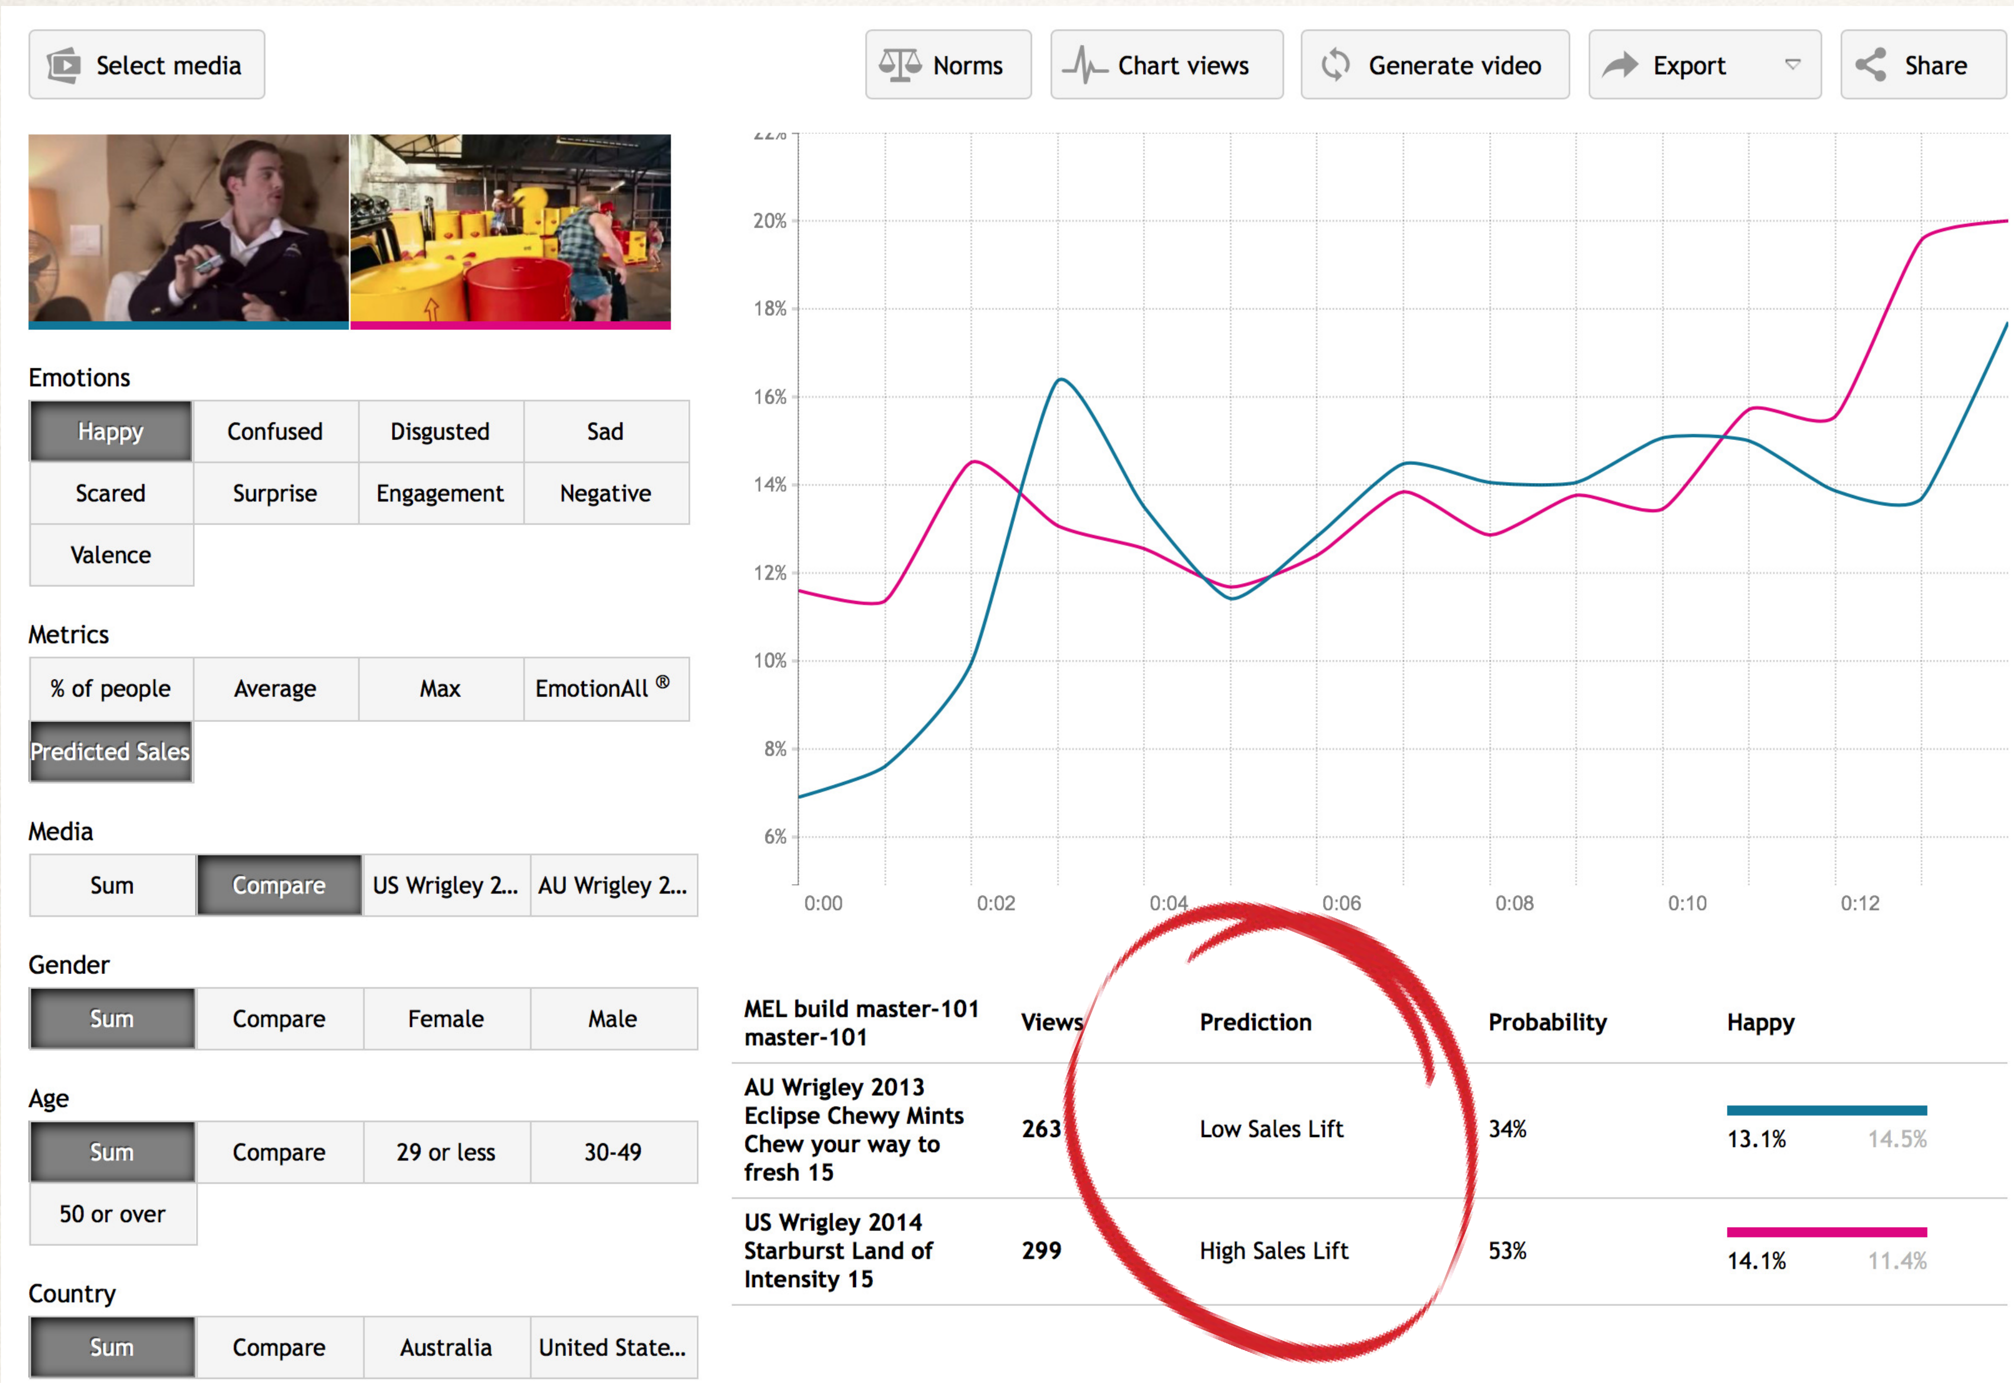

realeyes

Step #2: Find optimal viewer segmentation (M14-M16)

- Challenge 1: find robust segmentation method for panels of moderate size
- Implement, test and validate **score driven** brute force **segmentation**
- Implement, test and validate similarity based **clustering** using emotion response and/or meta-info based viewer representation
- Idea: cluster responses based on **deviation from average**

Step #3: Recommend ads to most similar target audience (M16-M18)

- **Challenge 2**: define representations and similarity metric that can be used to identify target audience
- **Target variables**: Prediction accuracy, Brand Lift scores, View Through Conversion Quartiles, Click Through Rates
- **Benchmarking**: External validation with our partners

Work from M19

- Expand and improve ad recommendation engine with new **sentiment analysis technology** provided by SEWA partners. Evaluation will be done against our own baseline solution
- Improve audience segmentation by using non-sales data like **social media activity** (#views, likes or shares on Youtube, Facebook, Twitter)
- "**Instant prediction**" of emotion responses from audio-visual content and context. This would yield approximate ad matching and **fast targeting**.

Objectives

- Report on user requirements for each SEWA application (D7.1)
- Initial version of the Ad Recommendation Engine (D7.2)
- Second version of the Ad Recommendation Engine (D7.4)
- Final version of the sentiment-driven Ad Recommendation Engine (D7.6)

WP7:
Ad Recommendation Engine

Realeyes
Elnar Hajiyev, CTO
Gabor Szirtes, Head Scientist
Pierre Lupi Chen, Data Scientist

sewa
Automatic Sentiment Analysis in the Wild