Steering User Behavior with Badges

Yichen Yang (yy18), Lee Wang (lxwang2), Suneer Angra (angra2)

Ashton Anderson
Stanford University

Daniel Huttenlocher
Cornell University

Jon Kleinberg
Cornell University

Jure Leskovec
Stanford University

Related works

- User Participation and contribution in online domain [1,4].
- Badges a part of growing phenomenon of gamification [2].
- Credentialing system and Incentive system [3, 5].

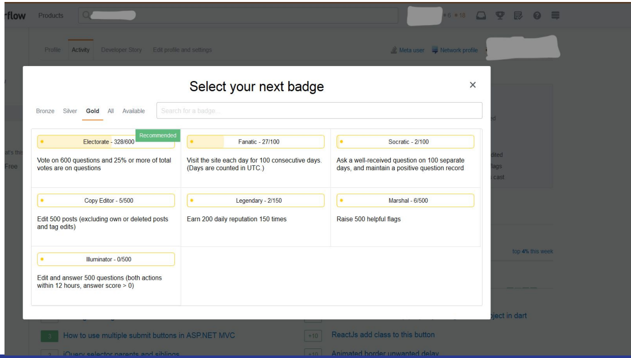

Introduction

- **Motivation**: Steer user behaviour to different kind of activities.
- **Contribution**: Propose a formal framework for reasoning about the effects badges will have, which can potentially be used for badge design.
- **Structure**:
    a. *Theoretical Model of User Behaviour* - for the user's optimization problem in the case of an arbitrary monotone badge
    b. *Empirical Evaluation* - Validate the model based on real data from Stack Overflow
    c. *Badge Placement* - Use the model to define badges with the goal of achieving a desired pattern of behavior

Assumptions and Limitations

In this model, the author only consider the badges with that are awarded once the user has reached a certain level of cumulative contribution.

The setup of the model

Here are the parameters:

1. Action types: $A = \{A_1, A_2, \ldots, A_n, A_{n+1}\}$, each A_i s one action type and the final one A_{n+1} is the life-action.
2. User histories: $\mathbf{a} \in \mathbb{R}^{n+1}$, each coordinate representing the accumulated number of the action and $\mathbf{a}^{n+1}$ is the number of life-action.
3. Unit user action: $\mathbf{e}_i$, with 1 coordinate has value = 1, and 0 elsewhere.
4. Badges and boundaries: *b*, $I_b(\mathbf{a})$, if $I_b(\mathbf{a})$ = 1, it means the the user history **a** warrants the badge *b*. If $I_b(\mathbf{a})$ = 1 and $I_b(\mathbf{a} - \mathbf{e}_i) = 0$ for some unit vector $\mathbf{e}_i$ it is the boundary of badge *b*.

The setup of the model

Here are the parameters:

5. Utilities and incentives:
   a. "Ideal" distribution **p**: user preferred action distribution without affecting by badges. User will sample from this distribution to determine next action.
   b. The cost of choosing another distribution **x** from **p**: $g(\mathbf{x}, \mathbf{p})$, the higher the deviation of a user from his preferred actions, the higher the cost g(·), here the author use $g(\mathbf{x}, \mathbf{p}) = \|\mathbf{x} - \mathbf{p}\|_2^2$ to calculate the cost of choosing different distribution. If **x** == **p**, $g(\mathbf{x}, \mathbf{p})$ = 0.
   c. The utility of achieving a badge: V_b, for each $b \in B$, we will have a corresponding V_b showing the utility of achieving the badge.

The setup of the model

Here are the parameters:

6. Exogenous probability: $\delta > 0$, after every action, user has probability = δ permanently leaves the system. So we only have $\theta = 1 - \delta$ probability that the user survives to perform next step.
7. User's policy: We call this choice of distributions $\mathcal{X} = \{\mathbf{x_a}\}$ the user's policy and $U(\mathbf{x_a})$ the utility of user in state **a** receives from the policy.

$$U(\mathbf{x_a}) = \sum_{b \in B} I_b(\mathbf{a}) V_b + \theta \sum_{i=1}^{n+1} \mathbf{x_a}^i \cdot U(\mathbf{x_{a+e_i}}) - g(\mathbf{x_a}, \mathbf{p})$$

The User's Optimization Problem

- How a user will behave under the model in the presence of a set of badges B
- The user chooses a policy $\mathcal{X} = \{\mathbf{x_a}\}$ to maximize his utility $U(\mathbf{x_0})$ starting from the origin
- Solved by the optimum of a Markov Decision Process (MDP)
    - But MDP is computationally expensive
- Developed an efficient algorithm
    - But it requires making use of the inherent structure of the problem as it arises from our model, rather than invoking a general class of results.
    - Focus on threshold badges that target one or two dimensions.
    - Can be extended to general badges.

One targeted dimension

- Here n = 2, i.e. 2 on-site action + life-action, 1 badge achieved by taking k action of type A_1 on-site action.
- Two key observations:

1. After the badge is achieved, no more utility can be gained, so the following behavior $\mathbf{x_a}$ will exactly as **p**.

2. If $\mathbf{a}$ and $\tilde{\mathbf{a}}$ have the same coordinate in dimension 1, then a sequence of actions starting at $\mathbf{a}$ crosses the badge boundary if and only if the same sequence starting at $\tilde{\mathbf{a}}$ crosses the badge boundary.

One targeted dimension

Here we just need to solve the utility function $U(\cdot)$ of $\mathbf{a}^1$, the number of A_1 actions the user has taken.

$$U(\mathbf{a}^1) = \theta \sum_{j=1}^{3} \mathbf{x}_{\mathbf{a}}^{j} \cdot U(\mathbf{x}_{\mathbf{a}+\mathbf{e}_j}) - g(\mathbf{x}_{\mathbf{a}}, \mathbf{p})$$

$$= \theta \cdot [\mathbf{x}_{\mathbf{a}}^{1} U(\mathbf{x}_{\mathbf{a}+\mathbf{e}_1}) + \mathbf{x}_{\mathbf{a}}^{2} U(\mathbf{x}_{\mathbf{a}}) + \mathbf{x}_{\mathbf{a}}^{3} U(\mathbf{x}_{\mathbf{a}})] - g(\mathbf{x}_{\mathbf{a}}, \mathbf{p})$$

then solving for $U(\mathbf{a}^1) = U(\mathbf{x}_{\mathbf{a}})$ we have

$$U(\mathbf{a}^1) = \frac{\theta \cdot \mathbf{x}_{\mathbf{a}}^{1} \cdot U(\mathbf{x}_{\mathbf{a}+\mathbf{e}_1}) - g(\mathbf{x}_{\mathbf{a}}, \mathbf{p})}{1 - \theta(\mathbf{x}_{\mathbf{a}}^{2} + \mathbf{x}_{\mathbf{a}}^{3})}$$

One targeted dimension

Since we have already computed $U(\mathbf{a}^1 + 1) = U(\mathbf{x}_{\mathbf{a}+\mathbf{e}_1})$, this becomes an optimization problem in 3 variables:

$$\underset{\mathbf{x}_\mathbf{a}}{\text{maximize}} \quad \frac{\theta \cdot \mathbf{x}_\mathbf{a}^1 \cdot C - g(\mathbf{x}_\mathbf{a}, \mathbf{p})}{1 - \theta(\mathbf{x}_\mathbf{a}^2 + \mathbf{x}_\mathbf{a}^3)}$$

$$\text{subject to} \quad \mathbf{x}_\mathbf{a}^j \geq 0, \ j = 1, 2, 3 \ \text{ and } \ \sum_{j=1}^{3} \mathbf{x}_\mathbf{a}^j = 1$$

where we've replaced $U(\mathbf{x}_{\mathbf{a}+\mathbf{e}_1})$ with C.

One targeted dimension

How to solve this problem:

Since we know $\mathbf{x_a} = \mathbf{p}$ for all states a such that $\mathbf{a}^1 \geq k$, we can use this to compute the optimal $\mathbf{x_a}$ for all a such that $\mathbf{a}^1 = k - 1$, and recurse all the way back to $\mathbf{a_0}$.

One targeted dimension Result

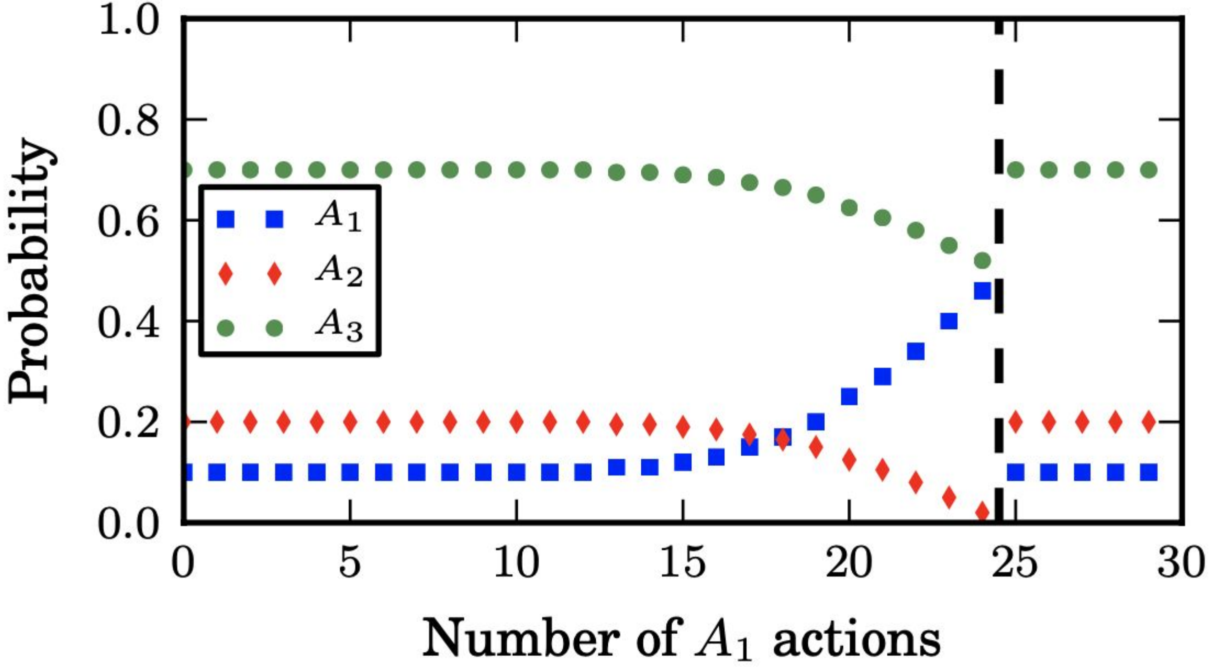

Extension to Multiple badges

multiple badges that all target the same dimension

We just need to consider $k_1 < k_2 < \ldots < k_m$

The value of each state is “initialized” with $\sum_{b \in B} I_b(\mathbf{a}) V_b$ and again our dynamic programming base case is that in all states a after all the badge boundaries, the user will choose $\mathbf{x_a} = \mathbf{p}$.

In general, the region between badges j – 1 and j is identical to the single badge case with a badge of value $V_{b_j} + U(\mathbf{x}_{k_j})$

Two targeted dimensions

Now we consider the case where different badges target different types of actions. We start with $B = \{b_1 = (k_1, 1), b_2 = (k_2, 2)\}$, so there are two dimensions with one badge targeting each (here let n = 2 for convenience).

Badge boundaries $\mathbf{a}^1 = k_1$ and $\mathbf{a}^2 = k_2$ split the action space into four regions:

- R: a finite rectangle bounded by the origin and $(k_1 - 1, k_2 - 1)$,
- H: an infinite horizontal strip with boundary points $(k_1, 0)$ and $(k_1, k_2 - 1)$ extending rightward,
- V: an infinite vertical strip with boundary points $(0, k_2)$ and $(k_1 - 1, k_2)$ extending upward, and
- Q: a quadrant rooted at (k_1, k_2).

Two targeted dimensions

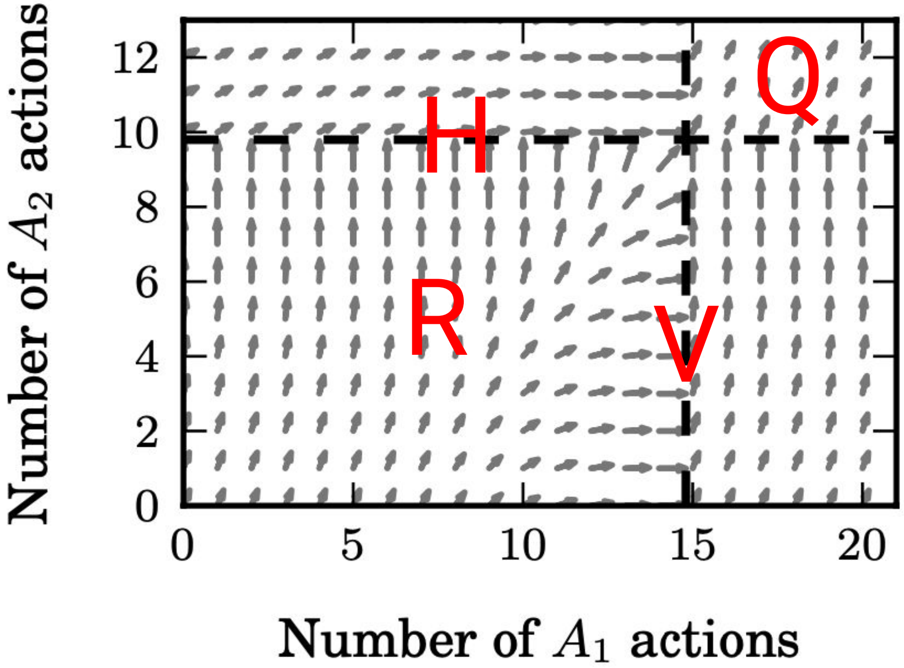

Two targeted dimensions

- Similarly to before, past all the badge boundaries ($\mathbf{a}^1 \geq k_1$ and $\mathbf{a}^2 \geq k_2$), the user has no incentive to deviate from **p**, so $\mathbf{x_a} = \mathbf{p}$ for all states in quadrant *Q*.
- Quadrants *H* and *V* are identical to the case of one threshold badge in one targeted dimension that we solved above in one targeted dimension.
- For the finite rectangle *R*, we can directly fill in in order of decreasing coordinate sum since the cells furthest from the origin depend on the value of states we already know from solving quadrants *Q*, *H* and *V*.

Two targeted dimensions

$$U(\mathbf{x_a}) = \theta \sum_{j=1}^{n+1} \mathbf{x_a}^j \cdot U(\mathbf{x_{a+e_j}}) - g(\mathbf{x_a}, \mathbf{p})$$

Consider a state $\mathbf{a}$ in region R that we process in order. We have already computed $U(\mathbf{x_{a+e_1}})$ and $U(\mathbf{x_{a+e_2}})$, so we can further simplify:

$$U(\mathbf{x_a}) = \frac{\theta \cdot (C_1 \cdot \mathbf{x_a}^1 + C_2 \cdot \mathbf{x_a}^2) - g(\mathbf{x_a}, \mathbf{p})}{1 - \theta \cdot \mathbf{x_a}^3}$$

where $C_j = U(\mathbf{x_{a+e_j}})$ for $j = 1, 2$ are the constants we have computed. This results in an optimization problem very similar to the one we had in the one-dimensional case:

Two targeted dimensions

$$\underset{\mathbf{x_a}}{\text{maximize}} \quad \frac{\theta \cdot (C_1 \cdot \mathbf{x_a}^1 + C_2 \cdot \mathbf{x_a}^2) - g(\mathbf{x_a}, \mathbf{p})}{1 - \theta \cdot \mathbf{x_a}^3}$$

$$\text{subject to} \quad \mathbf{x_a}^j \geq 0, \; j = 1, 2, 3 \;\; \text{and} \;\; \sum_{j=1}^{3} \mathbf{x_a}^j = 1$$

Two targeted dimensions

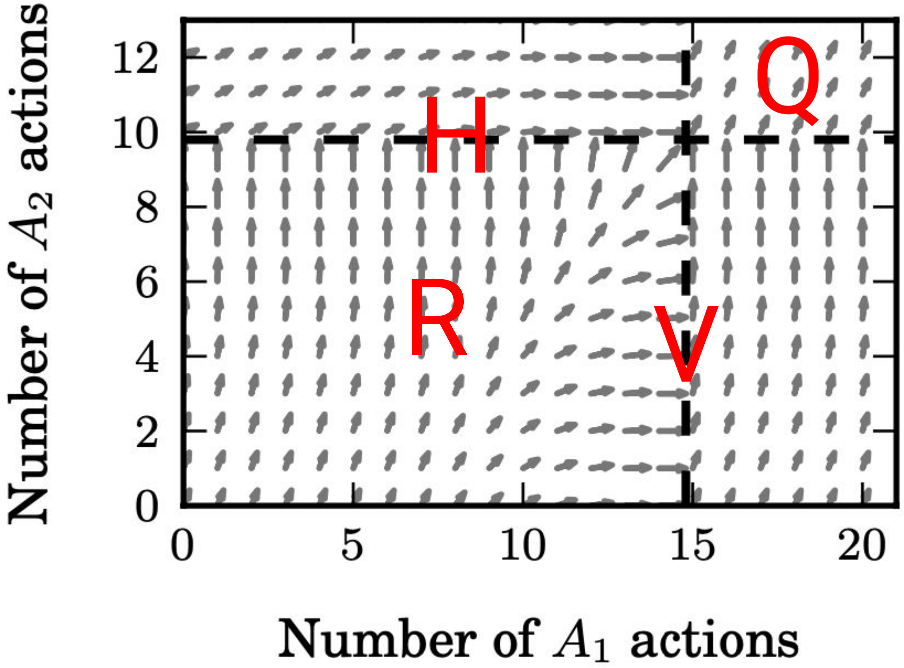

General Monotone Badges

Problem statement: There is an arbitrary monotone badge.

Dickson's Lemma: any monotone subset of $\mathbb{N}^m$ has only finitely many minimal elements. It shows that obtain badge *b* necessarily have finite description.

Key fact: any action sequence starting from $\mathbf{a}$ leads to obtaining the badge if and only if the same action sequence starting from $\tilde{\mathbf{a}}$ does.

Therefore, we can set $\mathbf{x}_{\mathbf{a}} = \mathbf{x}_{\tilde{\mathbf{a}}}$ in the optimal policy and compute the optimal policy on the subset.

Empirical Evaluation on Stack Overflow

- How do the predictions of the model compare with the aggregate behavior of individuals on Stack Overflow?
- To what extent do badges steer user behavior?
- What insights does the model provide to site designers who want to maximize the effectiveness of badges?

Evaluated Stack Overflow Badges

- Over 100 different badges on Stack Overflow, but only 2 badges are evaluated
  - "Electorate Badge": awarded for taking at least 600 Q-votes and having at least one Q-vote for every four A-votes
  - "Civic Duty Badge": awarded after voting 300 times (on questions or answers)
- Exclude one shot badges (e.g. "Great Question" badge for a single high-quality contribution)

Badge Details

- Assume that users can see their progress towards the badge and how many more actions they need
- Also, only considering badges with a prespecified number of actions (600 for "Electorate" badge and 300 for "Civic Duty" badge)

Collected Stack Overflow data

- Data from site's inception on 2008 to 2010
- Each individual action performed by a user is recorded and timestamped
- Complete sequence of actions that users take and measure their progress towards obtaining badges can be observed

Activity Around Badge Boundary

- For each user, bin the number of actions of each type by day
- For each badge, take the complete set of users who ever obtained that badge and axis-align their activity profiles

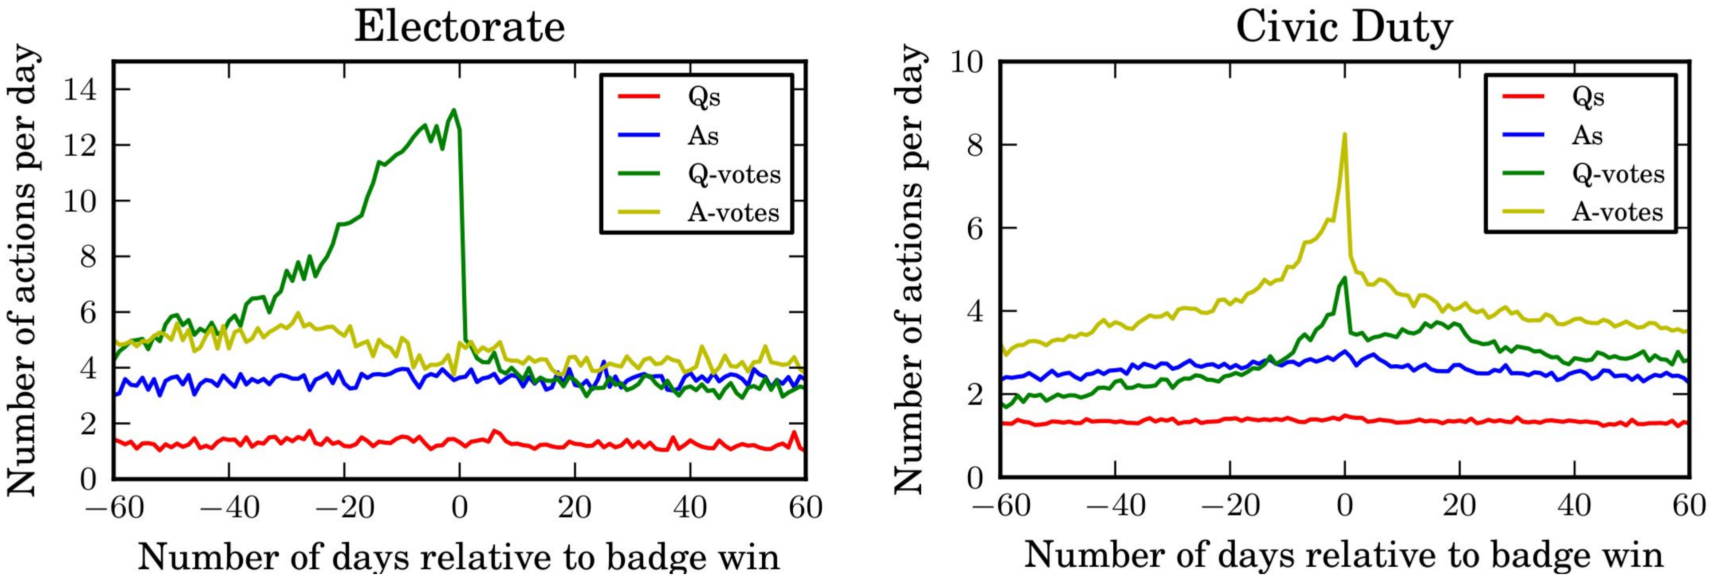

Figure 3: Number of actions per day as a function of number of days relative the time of obtaining a badge. Notice steering in the sense of increased activity on actions targeted by the badge.

Activity Around Badge Boundary Insights

- Clear steering towards badge boundary
- Acceleration effect when approaching the boundary
- Activity drops off to normal level after obtaining the badge
- Badges increased overall user participation on the Stack Overflow

Turning towards the badge

- Given the increase in user activity around the badge boundary, *how* do users steer towards the badge boundary?

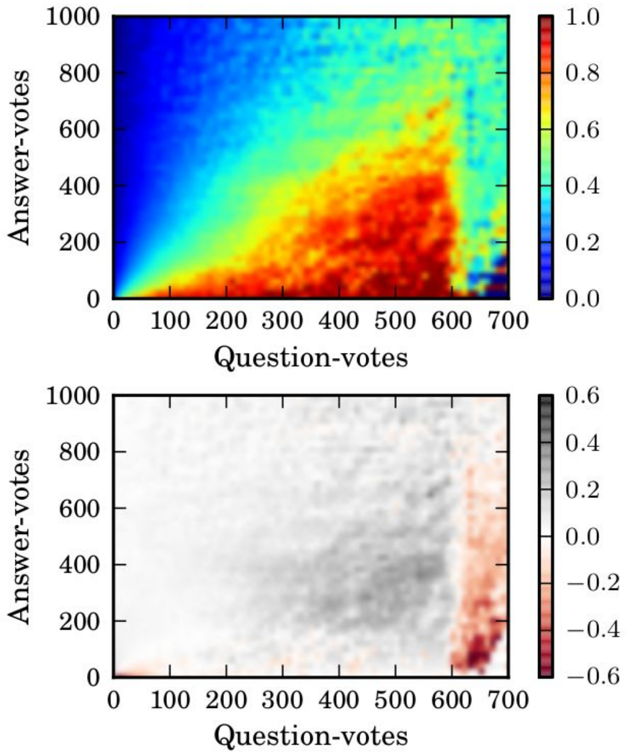

Figure 4: Electorate badge. Given that user has taken x question-votes and y answer-votes, what is the probability that next action will be a question vote. Top: Raw probability. Bottom: Relative change in probability of question-voting. Notice the effects of "turning" towards the badge boundary.

- Users shift their effort on the site towards actions that lead to badges

Designing Badges

- We as site designers want to maximize steering
  - How much user steering do different badge placements provide?
  - How might site designers place badges to best achieve desired user behavior?

Definition

- *Yield* is the total fraction of actions on a targeted action that results from a particular set of badge placements
- Higher yield means badge is more effective
- *Gain* is the difference between the yield and the default fraction of actions the user takes in the absence of badges

Setup - One Badge, One dimension

- Consider a site designer who wants to maximize the yield for a single badge on a single dimension (action type) called A_1
- Optimal badge placement requires balance between two competing forces
  - Threshold should be high enough so it effects many actions
  - However, if the threshold is too high then users won't live long enough on the site to achieve it so they won't be incentivized to steer towards it
- Define $\mathbf{p}^1$ as the user's preferred probability for taking A_1

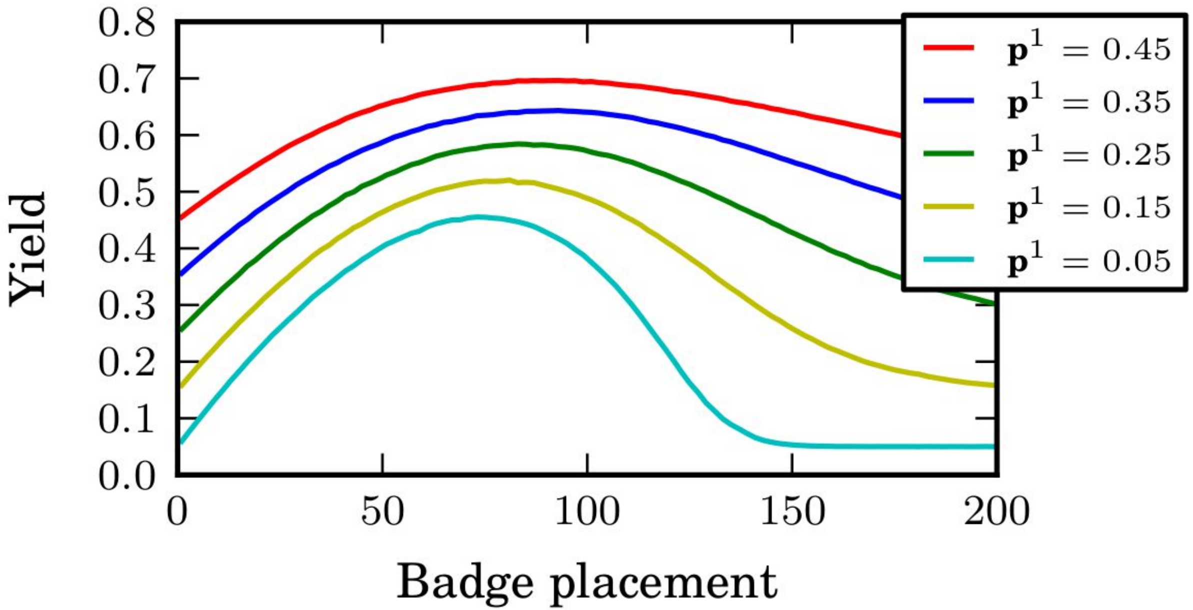

Figure 5: The resulting fraction of actions on the targeted dimension (here A_1) as a function of where the badge is placed. The different curves show how the relationship varies as the user's preferences p change.

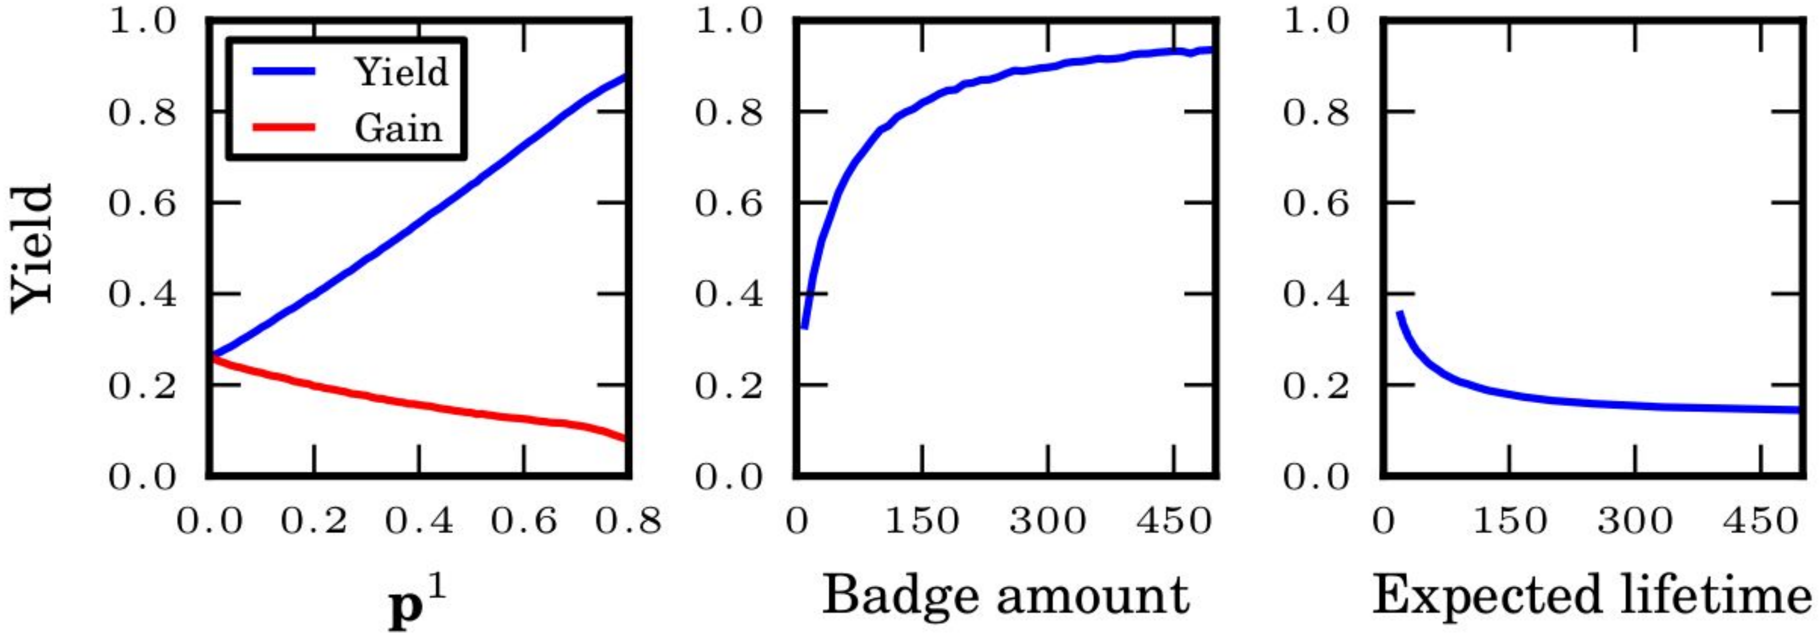

Figure 6: How optimal yield varies with setting parameters.

One Badge, One Dimension Insights

- The effectiveness of a badge is maximized at a surprisingly high internal optimum
  - $\mathbf{p}^1 = 0.05$ implies that the user would take only 5 A_1 actions in the absence of badges (assuming a total of 100 actions), yet the optimal badge location is far away at $A_1 = 75$
- Users are steered more on the actions they dislike (i.e. low $\mathbf{p}^1$) than those they like
- Optimal badge placement increases with $\mathbf{p}^1$
- Badges have stronger effects when users' lives on the site are shorter
- Diminishing returns when increasing the badge's value

Two badges, One Dimension

- How should a site designer place two badges that target the same dimension to maximize yield?
- What value (i.e. utility) should be assigned for the pair of badges to maximize yield?

Placement of two badges for maximum yield

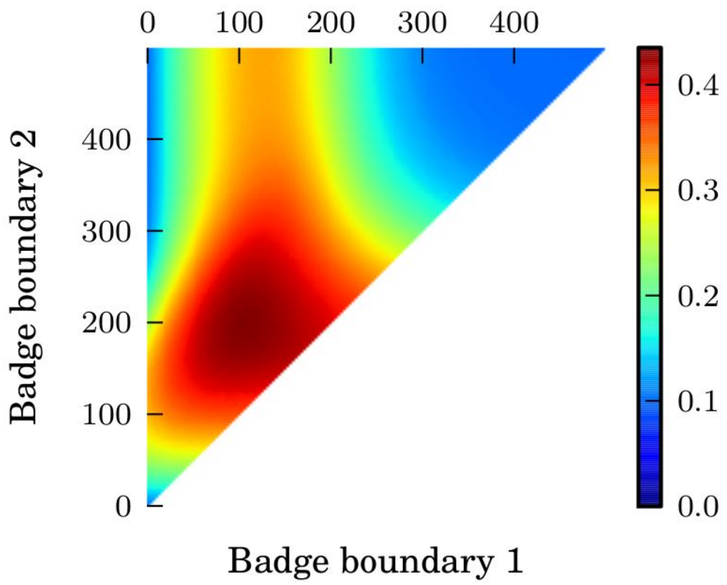

Figure 7: Yield as a function of where two badges are placed on the same targeted dimension. Every cell (x, y) corresponds to placing one badge at $a^1 = x$ actions and another at $a^1 = y$ actions, and the color of the cell indicates the resulting yield.

Note: Both badges have the same value

Placement of Two Badges Insights

- Better to place the two badges at distinct locations rather than combining them into a single large badge
- Badges that are spaced approximately evenly apart is optimal

Value for Two Badges to Maximize Yield

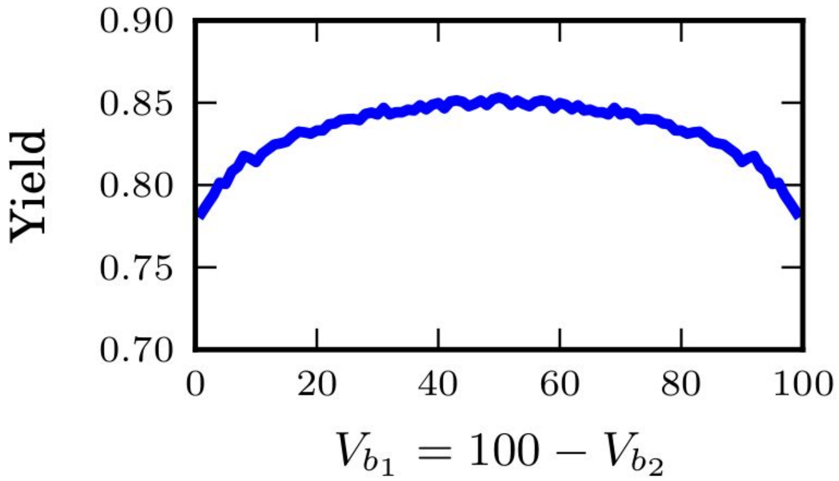

Figure 8: How the optimal yield achievable with two badges depends on how a fixed amount of utility is split across the two badges. The more even the split, the higher the yield.

Value for Two Badges Insight

- The more even the value split, the higher the yield
- Designer should create badges so that have about equal value

Two Badges, Two Dimensions Insights

- Three scenarios:
  - Put both on A_1
  - Put both on A_2
  - One on each
- Maximum yield on a single dimension occurs when placing both badges on that dimension
- High yield on both dimensions occur when placing one badge on each dimension
- Certain yields of A_1 and A_2 in conjunction are not possible using any combination of the two badges

Conclusion

- Badge can steer user behaviour and also increase their overall participation.
- Site designer can influence the value of the badge, which also depends on the structure of the possible actions, user's preferences for these actions, and user's expected lifetimes.

Strength and weakness

- Strengths:
  - Tries to formalize the framework for Badge Incentives
  - Provides valuable insights for site designers looking to optimize user participation
  - Concrete evidence of users changing their behavior in response to badges
- Weakness:
  - The model is mainly based on the badges for k repeated actions. However, it doesn't show the influence of badges for 1 action (a specific action). This can be extended by the future papers or other models.
  - Behavior ?? The model doesn't take into account the "quality" of user actions.
    - Incorporating action quality into the model could be future work
  - User could value badges differently depending on which other badges have already been obtained -- this could be incorporated into the value function V

Q&A

Thank you!

References

[1] H. L. O’Brien and E. G. Toms. What is user engagement? A conceptual framework for defining user engagement with Technology. *Journal of the American Society for Information Science and Technology*, 59(6):938–955, 2008.

[2] S. Deterding, M. Sicart, L. Nacke, K. O’Hara, and D. Dixon. Gamification: Using game-design elements in non-gaming contexts. In *CHI Workshop on Gamification*, 2011.

[3] J. H. Bishop. Incentives for learning: Why american high school students compare so poorly to their counterparts overseas. Technical Report CAHRS 89-09, Cornell University School of Industrial and Labor Relations, 1989.

[4] N. Singer. You’ve won a badge (and now we know all about you). *New York Times*, 4 February 2012.

[5] J. Antin and E. Churchill. Badges in social media: A social psychological perspective. In *CHI Workshop on Gamification*, 2011.