Identifying platform effects in social media data

Momin M. Malik[1] <momin.malik@cs.cmu.edu>
Jürgen Pfeffer[1,2] <juergen.pfeffer@tum.de>

[1] Institute for Software Research
School of Computer Science
Carnegie Mellon University

[2] Bavarian School of Public Policy
Technical University of Munich

ICWSM

Cologne, Germany
May 18, 2016

Carnegie Mellon University School of Computer Science

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., "Throwback Thursday," #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Carnegie Mellon University
School of Computer Science

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., “Throwback Thursday,” #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., "Throwback Thursday," #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., "Throwback Thursday," #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., "Throwback Thursday," #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Carnegie Mellon University School of Computer Science

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., “Throwback Thursday,” #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., "Throwback Thursday," #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., "Throwback Thursday," #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., "Throwback Thursday," #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Biases in social media data

Multiple forms of bias (Tufekci 2014; Ruths and Pfeffer 2014):

- Demographics (e.g., Reddit vs Instagram users)
- Heterogeneous users (e.g., corporate users, celebrity fans, activists)
- Culture and norms (e.g., "Throwback Thursday," #tbt)
- National differences (Problete et al. 2011)
- APIs sample nonrandomly (Morstatter et al. 2014)
- Changes over time (Liu et al. 2014)

Why do we care? Because these biases are why research findings fail to generalize across platforms, time, and users (e.g., Cohen and Ruths 2013)

Carnegie Mellon University School of Computer Science

Motivation
Data
Method
Findings
Conclusions

An unaddressed bias: "platform effects"

Platform effects are the ways in which the *design* and *technical features* of a given platform *constrain*, *distort*, and *shape* user behavior.

- Simple example: 140 character limit is a hard constraint on human behavior
- We know companies try very hard to shape user behavior in certain ways (van Dijck 2013; Gehl 2014)
- Are we studying behavior? Or just the successful algorithmic management of users?

An unaddressed bias: “platform effects”

Platform effects are the ways in which the *design* and *technical features* of a given platform *constrain*, *distort*, and *shape* user behavior.

- Simple example: 140 character limit is a hard constraint on human behavior
- We know companies try very hard to shape user behavior in certain ways (van Dijck 2013; Gehl 2014)
- Are we studying behavior? Or just the successful algorithmic management of users?

An unaddressed bias: "platform effects"

Platform effects are the ways in which the *design* and *technical features* of a given platform *constrain*, *distort*, and *shape* user behavior.

- Simple example: 140 character limit is a hard constraint on human behavior

- We know companies try very hard to shape user behavior in certain ways (van Dijck 2013; Gehl 2014)

- Are we studying behavior? Or just the successful algorithmic management of users?

An unaddressed bias: "platform effects"

Platform effects are the ways in which the *design* and *technical features* of a given platform *constrain*, *distort*, and *shape* user behavior.

- Simple example: 140 character limit is a hard constraint on human behavior
- We know companies try very hard to shape user behavior in certain ways (van Dijck 2013; Gehl 2014)
- Are we studying behavior? Or just the successful algorithmic management of users?

An unaddressed bias: “platform effects”

Platform effects are the ways in which the *design* and *technical features* of a given platform *constrain*, *distort*, and *shape* user behavior.

- Simple example: 140 character limit is a hard constraint on human behavior
- We know companies try very hard to shape user behavior in certain ways (van Dijck 2013; Gehl 2014)
- Are we studying behavior? Or just the successful algorithmic management of users?

Carnegie Mellon University School of Computer Science

Motivation

Data

Method

Findings

Conclusions

Studying platform effects

- How can we know about platform effects?
  - A/B tests from inside companies? Usually not accessible
  - Reverse engineering from inputs and outputs (Diakopoulous 2014)? Not quite appropriate here
- Instead: find *natural experiments*.
- One great place to look: what we think of as "data artifacts," discontinuities in data because of some system change
- The discontinuities tell us about the causal impact of the engineering changes!

Carnegie Mellon University School of Computer Science

Motivation
Data
Method
Findings
Conclusions

Studying platform effects

- How can we know about platform effects?
  - A/B tests from inside companies? Usually not accessible
  - Reverse engineering from inputs and outputs (Diakopoulos 2014)? Not quite appropriate here
- Instead: find *natural experiments*.
- One great place to look: what we think of as "data artifacts," discontinuities in data because of some system change
- The discontinuities tell us about the causal impact of the engineering changes!

Carnegie Mellon University School of Computer Science

Studying platform effects

- How can we know about platform effects?
  - A/B tests from inside companies? Usually not accessible
  - Reverse engineering from inputs and outputs (Diakopoulos 2014)? Not quite appropriate here
- Instead: find *natural experiments*.
- One great place to look: what we think of as "data artifacts," discontinuities in data because of some system change
- The discontinuities tell us about the causal impact of the engineering changes!

Studying platform effects

- How can we know about platform effects?
  - A/B tests from inside companies? Usually not accessible
  - Reverse engineering from inputs and outputs (Diakopoulous 2014)? Not quite appropriate here
- Instead: find *natural experiments*.
- One great place to look: what we think of as "data artifacts," discontinuities in data because of some system change
- The discontinuities tell us about the causal impact of the engineering changes!

Studying platform effects

- How can we know about platform effects?
  - A/B tests from inside companies? Usually not accessible
  - Reverse engineering from inputs and outputs (Diakopoulous 2014)? Not quite appropriate here
- Instead: find *natural experiments*.
- One great place to look: what we think of as "data artifacts," discontinuities in data because of some system change
- The discontinuities tell us about the causal impact of the engineering changes!

Studying platform effects

- How can we know about platform effects?
  - A/B tests from inside companies? Usually not accessible
  - Reverse engineering from inputs and outputs (Diakopoulous 2014)? Not quite appropriate here
- Instead: find *natural experiments*.
- One great place to look: what we think of as "data artifacts," discontinuities in data because of some system change
- The discontinuities tell us about the causal impact of the engineering changes!

Studying platform effects

- How can we know about platform effects?
  - A/B tests from inside companies? Usually not accessible
  - Reverse engineering from inputs and outputs (Diakopoulous 2014)? Not quite appropriate here
- Instead: find *natural experiments*.
- One great place to look: what we think of as "data artifacts," discontinuities in data because of some system change
- The discontinuities tell us about the causal impact of the engineering changes!

Carnegie Mellon University School of Computer Science

Two cases

- Netflix Prize data (see Koren 2009)
- Facebook New Orleans data (Zignani et al. 2014, via Viswanath et al. 2011)

Two cases

- Netflix Prize data (see Koren 2009)
- Facebook New Orleans data (Zignani et al. 2014, via Viswanath et al. 2011)

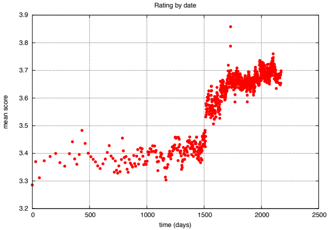

Fig. 1 from Koren (2009): Temporal effects emerging within the Netflix movie rating dataset. The average movie rating made a sudden jump in early 2004 (1500 days since the first rating in the dataset). Each point averages 100,000 rating instances.

Two cases

- Netflix Prize data (see Koren 2009)
- Facebook New Orleans data (Zignani et al. 2014, via Viswanath et al. 2011)

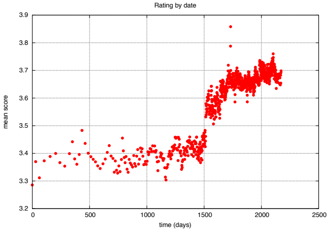

Fig. 1 from Koren (2009): Temporal effects emerging within the Netflix movie rating dataset. The average movie rating made a sudden jump in early 2004 (1500 days since the first rating in the dataset). Each point averages 100,000 rating instances.

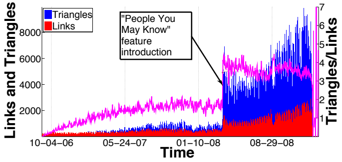

(b) Facebook New Orleans

Fig. 2 from Zignani et al. (2014): Number of new links (red) and triangles (blue) formed during the growth of Facebook New Orleans, sampled each day. The magenta line represents the ratio between the triangle and the links created in a day (y-scale on the right).

Carnegie Mellon University School of Computer Science

Causal framework

- $Y_i|set(T=1)$ is value of Y_i if i is given treatment T, $Y_i|set(T=0)$ is value of Y_i if i is not given treatment T
- For a given i, can never observe both
- Instead, use expectations. Define the *average treatment effect* α as

$$\alpha := E(Y_i|set(T=1)) - E(Y_i|set(T=0))$$

- If $(Y_i|set(T=1), Y_i|set(T=0))$ is independent of of T (what randomization does), then

$$E(Y_i|set(T=1)) = E(Y_i|T=1) \text{ and } E(Y_i|set(T=0)) = E(Y_i|T=0)$$

Carnegie Mellon University School of Computer Science

Causal framework

- $Y_i|set(T=1)$ is value of Y_i if i is given treatment T, $Y_i|set(T=0)$ is value of Y_i if i is not given treatment T
- For a given i, can never observe both
- Instead, use expectations. Define the *average treatment effect* α as

$$\alpha := E(Y_i|set(T=1)) - E(Y_i|set(T=0))$$

- If $(Y_i|set(T=1), Y_i|set(T=0))$ is independent of of T (what randomization does), then

$$E(Y_i|set(T=1)) = E(Y_i|T=1) \text{ and } E(Y_i|set(T=0)) = E(Y_i|T=0)$$

Carnegie Mellon University School of Computer Science

Causal framework

- $Y_i|set(T=1)$ is value of Y_i if i is given treatment T, $Y_i|set(T=0)$ is value of Y_i if i is not given treatment T
- For a given i, can never observe both
- Instead, use expectations. Define the *average treatment effect* α as

$$\alpha := E(Y_i|set(T=1)) - E(Y_i|set(T=0))$$

- If $(Y_i|set(T=1), Y_i|set(T=0))$ is independent of of T (what randomization does), then

$$E(Y_i|set(T=1)) = E(Y_i|T=1) \text{ and } E(Y_i|set(T=0)) = E(Y_i|T=0)$$

Causal framework

- $Y_i | set(T = 1)$ is value of Y_i if i is given treatment T, $Y_i | set(T = 0)$ is value of Y_i if i is not given treatment T
- For a given i, can never observe both
- Instead, use expectations. Define the *average treatment effect* α as

$$\alpha := E(Y_i | set(T = 1)) - E(Y_i | set(T = 0))$$

- If $(Y_i | set(T = 1), Y_i | set(T = 0))$ is independent of of T (what randomization does), then

$$E(Y_i | set(T = 1)) = E(Y_i | T = 1) \text{ and } E(Y_i | set(T = 0)) = E(Y_i | T = 0)$$

Causal framework

- $Y_i|set(T=1)$ is value of Y_i if i is given treatment T, $Y_i|set(T=0)$ is value of Y_i if i is not given treatment T
- For a given i, can never observe both
- Instead, use expectations. Define the *average treatment effect* α as

$$\alpha := E(Y_i|set(T=1)) - E(Y_i|set(T=0))$$

- If $(Y_i|set(T=1), Y_i|set(T=0))$ is independent of of T (what randomization does), then

$$E(Y_i|set(T=1)) = E(Y_i|T=1) \text{ and } E(Y_i|set(T=0)) = E(Y_i|T=0)$$

Regression discontinuity

Regression Discontinuity (RD) Design (Imbens and Lemieux 2008) is the use of a treatment that is effective strictly above some cutoff value c of a covariate X_i, $T = \mathbf{1}(X_i > c)$.

We can make a point estimate of the effect of treatment on the treated, which is the *local average treatment effect* (Imbens and Angrist 1994). Then,

$$\alpha = E(Y_i|\mathbf{1}(X_i > c) = 1) - E(Y_i|\mathbf{1}(X_i > c) = 0)$$

$$= \lim_{x \downarrow c} E(Y_i|X_i = x) - \lim_{x \uparrow c} E(Y_i|X_i = x)$$

Regression discontinuity

Regression Discontinuity (RD) Design (Imbens and Lemieux 2008) is the use of a treatment that is effective strictly above some cutoff value c of a covariate X_i, $T = \mathbf{1}(X_i > c)$.

We can make a point estimate of the effect of treatment on the treated, which is the *local average treatment effect* (Imbens and Angrist 1994). Then,

$$\alpha = E(Y_i|\mathbf{1}(X_i > c) = 1) - E(Y_i|\mathbf{1}(X_i > c) = 0)$$

$$= \lim_{x \downarrow c} E(Y_i|X_i = x) - \lim_{x \uparrow c} E(Y_i|X_i = x)$$

Regression discontinuity

Regression Discontinuity (RD) Design (Imbens and Lemieux 2008) is the use of a treatment that is effective strictly above some cutoff value c of a covariate X_i, $T = \mathbf{1}(X_i > c)$.

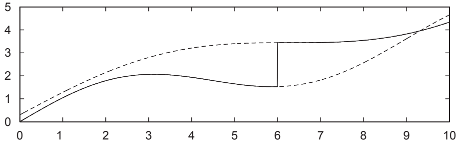

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

We can make a point estimate of the effect of treatment on the treated, which is the *local average treatment effect* (Imbens and Angrist 1994). Then,

$$\alpha = E(Y_i|\mathbf{1}(X_i > c) = 1) - E(Y_i|\mathbf{1}(X_i > c) = 0)$$

$$= \lim_{x \downarrow c} E(Y_i|X_i = x) - \lim_{x \uparrow c} E(Y_i|X_i = x)$$

Regression discontinuity

Regression Discontinuity (RD) Design (Imbens and Lemieux 2008) is the use of a treatment that is effective strictly above some cutoff value c of a covariate X_i, $T = \mathbf{1}(X_i > c)$.

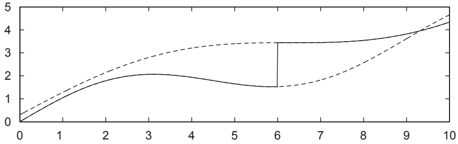

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

We can make a point estimate of the effect of treatment on the treated, which is the *local average treatment effect* (Imbens and Angrist 1994). Then,

$$\alpha = E(Y_i | \mathbf{1}(X_i > c) = 1) - E(Y_i | \mathbf{1}(X_i > c) = 0)$$

$$= \lim_{x \downarrow c} E(Y_i | X_i = x) - \lim_{x \uparrow c} E(Y_i | X_i = x)$$

Regression discontinuity

Regression Discontinuity (RD) Design (Imbens and Lemieux 2008) is the use of a treatment that is effective strictly above some cutoff value c of a covariate X_i, $T = \mathbf{1}(X_i > c)$.

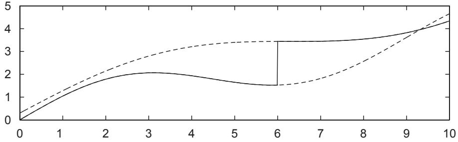

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

We can make a point estimate of the effect of treatment on the treated, which is the *local average treatment effect* (Imbens and Angrist 1994). Then,

$$\alpha = E(Y_i|\mathbf{1}(X_i > c) = 1) - E(Y_i|\mathbf{1}(X_i > c) = 0)$$

$$= \lim_{x \downarrow c} E(Y_i|X_i = x) - \lim_{x \uparrow c} E(Y_i|X_i = x)$$

Regression discontinuity

Regression Discontinuity (RD) Design (Imbens and Lemieux 2008) is the use of a treatment that is effective strictly above some cutoff value c of a covariate X_i, $T = \mathbf{1}(X_i > c)$.

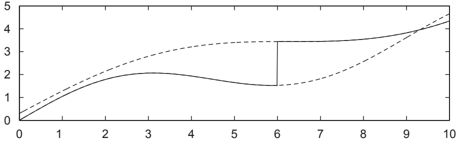

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

We can make a point estimate of the effect of treatment on the treated, which is the *local average treatment effect* (Imbens and Angrist 1994). Then,

$$\alpha = E(Y_i | \mathbf{1}(X_i > c) = 1) - E(Y_i | \mathbf{1}(X_i > c) = 0)$$

$$= \lim_{x \downarrow c} E(Y_i | X_i = x) - \lim_{x \uparrow c} E(Y_i | X_i = x)$$

Regression discontinuity

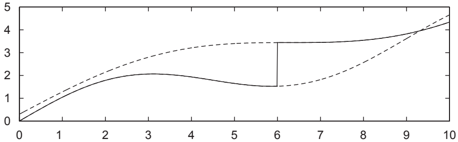

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

Linear univariate case:

Can generalize by doing separate nonparametric fits (e.g., local linear) on either side of the discontinuity

Regression discontinuity

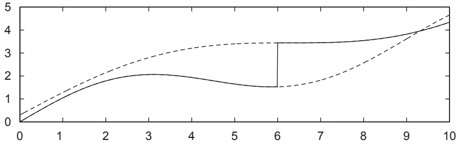

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

Linear univariate case:

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 \mathbf{1}(x_i > c) + \beta_3 x_i \mathbf{1}(x_i > c) + \varepsilon_i$$

$$\implies \widehat{\alpha} = \lim_{x \downarrow c} \widehat{E}(Y_i | X_i = x) - \lim_{x \uparrow c} \widehat{E}(Y_i | X_i = x)$$

$$= (\widehat{\beta_0} + \widehat{\beta_1} c + \widehat{\beta_2} + \widehat{\beta_3} c) - (\widehat{\beta_0} + \widehat{\beta_1} c) = \boxed{\widehat{\beta_2} + \widehat{\beta_3} c}$$

Can generalize by doing separate nonparametric fits (e.g., local linear) on either side of the discontinuity

Regression discontinuity

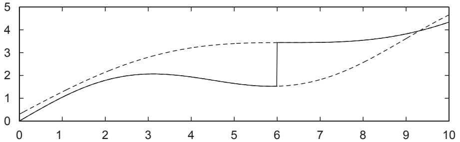

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

Linear univariate case:

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 \mathbf{1}(x_i > c) + \beta_3 x_i \mathbf{1}(x_i > c) + \varepsilon_i$$

$$\implies \widehat{\alpha} = \lim_{x \downarrow c} \widehat{E}(Y_i | X_i = x) - \lim_{x \uparrow c} \widehat{E}(Y_i | X_i = x)$$

$$= (\widehat{\beta_0} + \widehat{\beta_1} c + \widehat{\beta_2} + \widehat{\beta_3} c) - (\widehat{\beta_0} + \widehat{\beta_1} c) = \boxed{\widehat{\beta_2} + \widehat{\beta_3} c}$$

Can generalize by doing separate nonparametric fits (e.g., local linear) on either side of the discontinuity

Regression discontinuity

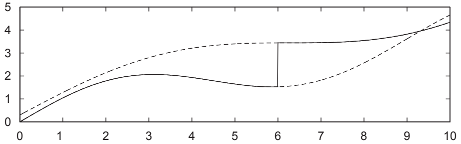

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

Linear univariate case:

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 \mathbf{1}(x_i > c) + \beta_3 x_i \mathbf{1}(x_i > c) + \varepsilon_i$$

$$\implies \widehat{\alpha} = \lim_{x \downarrow c} \widehat{E}(Y_i | X_i = x) - \lim_{x \uparrow c} \widehat{E}(Y_i | X_i = x)$$

$$= (\widehat{\beta_0} + \widehat{\beta_1} c + \widehat{\beta_2} + \widehat{\beta_3} c) - (\widehat{\beta_0} + \widehat{\beta_1} c) = \boxed{\widehat{\beta_2} + \widehat{\beta_3} c}$$

Can generalize by doing separate nonparametric fits (e.g., local linear) on either side of the discontinuity

Regression discontinuity

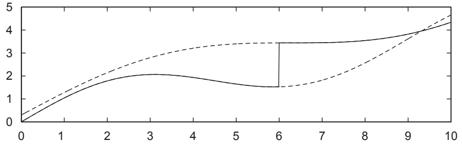

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

Linear univariate case:

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 \mathbf{1}(x_i > c) + \beta_3 x_i \mathbf{1}(x_i > c) + \varepsilon_i$$

$$\implies \widehat{\alpha} = \lim_{x \downarrow c} \widehat{E}(Y_i | X_i = x) - \lim_{x \uparrow c} \widehat{E}(Y_i | X_i = x)$$

$$= (\widehat{\beta_0} + \widehat{\beta_1} c + \widehat{\beta}_2 + \widehat{\beta}_3 c) - (\widehat{\beta_0} + \widehat{\beta_1} c) = \boxed{\widehat{\beta}_2 + \widehat{\beta}_3 c}$$

Can generalize by doing separate nonparametric fits (e.g., local linear) on either side of the discontinuity

Regression discontinuity

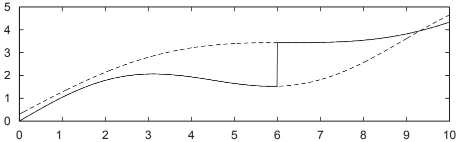

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

Linear univariate case:

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 \mathbf{1}(x_i > c) + \beta_3 x_i \mathbf{1}(x_i > c) + \varepsilon_i$$

$$\implies \widehat{\alpha} = \lim_{x \downarrow c} \widehat{E}(Y_i | X_i = x) - \lim_{x \uparrow c} \widehat{E}(Y_i | X_i = x)$$

$$= (\widehat{\beta_0} + \widehat{\beta_1} c + \widehat{\beta_2} + \widehat{\beta_3} c) - (\widehat{\beta_0} + \widehat{\beta_1} c) = \boxed{\widehat{\beta_2} + \widehat{\beta_3} c}$$

Can generalize by doing separate nonparametric fits (e.g., local linear) on either side of the discontinuity

Regression discontinuity

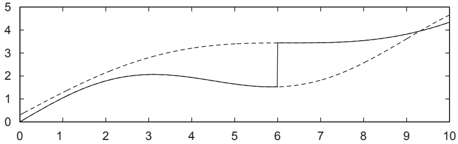

Fig. 2 from Imbens and Lemieux (2008): Potential and observed outcome regression functions.

Linear univariate case:

$$Y_i = \beta_0 + \beta_1 x_i + \beta_2 \mathbf{1}(x_i > c) + \beta_3 x_i \mathbf{1}(x_i > c) + \varepsilon_i$$

$$\implies \widehat{\alpha} = \lim_{x \downarrow c} \widehat{E}(Y_i | X_i = x) - \lim_{x \uparrow c} \widehat{E}(Y_i | X_i = x)$$

$$= (\widehat{\beta_0} + \widehat{\beta_1} c + \widehat{\beta}_2 + \widehat{\beta}_3 c) - (\widehat{\beta_0} + \widehat{\beta_1} c) = \boxed{\widehat{\beta}_2 + \widehat{\beta}_3 c}$$

Can generalize by doing separate nonparametric fits (e.g., local linear) on either side of the discontinuity

Carnegie Mellon University
School of Computer Science

Specification testing

- Want to do specification testing (say, is this really a discontinuity?). See if, for other choices of a cutoff, the confidence intervals to the left and right overlap
- Problem: RD not made for time series. Not accounting for temporal autocorrelation makes confidence intervals too narrow
- Alternatives?
  - Event Analysis, Interrupted Time Series: no CIs
  - ARIMA models: Differencing destroys discontinuity
  - Gaussian Process regression: CIs still too narrow
- Solution: tolerance intervals (empirical coverage), which we fit with quantile regression. Captures irreducible variance of time series, gives wide enough intervals

Specification testing

- Want to do specification testing (say, is this really a discontinuity?). See if, for other choices of a cutoff, the confidence intervals to the left and right overlap

- Problem: RD not made for time series. Not accounting for temporal autocorrelation makes confidence intervals too narrow

- Alternatives?
  - Event Analysis, Interrupted Time Series: no CIs
  - ARIMA models: Differencing destroys discontinuity
  - Gaussian Process regression: CIs still too narrow

- Solution: tolerance intervals (empirical coverage), which we fit with quantile regression. Captures irreducible variance of time series, gives wide enough intervals

Specification testing

- Want to do specification testing (say, is this really a discontinuity?). See if, for other choices of a cutoff, the confidence intervals to the left and right overlap
- Problem: RD not made for time series. Not accounting for temporal autocorrelation makes confidence intervals too narrow
- Alternatives?
  - Event Analysis, Interrupted Time Series: no CIs
  - ARIMA models: Differencing destroys discontinuity
  - Gaussian Process regression: CIs still too narrow
- Solution: tolerance intervals (empirical coverage), which we fit with quantile regression. Captures irreducible variance of time series, gives wide enough intervals

Specification testing

- Want to do specification testing (say, is this really a discontinuity?). See if, for other choices of a cutoff, the confidence intervals to the left and right overlap
- Problem: RD not made for time series. Not accounting for temporal autocorrelation makes confidence intervals too narrow
- Alternatives?
  - Event Analysis, Interrupted Time Series: no CIs
  - ARIMA models: Differencing destroys discontinuity
  - Gaussian Process regression: CIs still too narrow
- Solution: tolerance intervals (empirical coverage), which we fit with quantile regression. Captures irreducible variance of time series, gives wide enough intervals

Carnegie Mellon University School of Computer Science

Specification testing

- Want to do specification testing (say, is this really a discontinuity?). See if, for other choices of a cutoff, the confidence intervals to the left and right overlap
- Problem: RD not made for time series. Not accounting for temporal autocorrelation makes confidence intervals too narrow
- Alternatives?
  - Event Analysis, Interrupted Time Series: no CIs
  - ARIMA models: Differencing destroys discontinuity
  - Gaussian Process regression: CIs still too narrow
- Solution: tolerance intervals (empirical coverage), which we fit with quantile regression. Captures irreducible variance of time series, gives wide enough intervals

Specification testing

- Want to do specification testing (say, is this really a discontinuity?). See if, for other choices of a cutoff, the confidence intervals to the left and right overlap
- Problem: RD not made for time series. Not accounting for temporal autocorrelation makes confidence intervals too narrow
- Alternatives?
  - Event Analysis, Interrupted Time Series: no CIs
  - ARIMA models: Differencing destroys discontinuity
  - Gaussian Process regression: CIs still too narrow
- Solution: tolerance intervals (empirical coverage), which we fit with quantile regression. Captures irreducible variance of time series, gives wide enough intervals

Specification testing

- Want to do specification testing (say, is this really a discontinuity?). See if, for other choices of a cutoff, the confidence intervals to the left and right overlap
- Problem: RD not made for time series. Not accounting for temporal autocorrelation makes confidence intervals too narrow
- Alternatives?
  - Event Analysis, Interrupted Time Series: no CIs
  - ARIMA models: Differencing destroys discontinuity
  - Gaussian Process regression: CIs still too narrow
- Solution: tolerance intervals (empirical coverage), which we fit with quantile regression. Captures irreducible variance of time series, gives wide enough intervals

Specification testing

- Want to do specification testing (say, is this really a discontinuity?). See if, for other choices of a cutoff, the confidence intervals to the left and right overlap
- Problem: RD not made for time series. Not accounting for temporal autocorrelation makes confidence intervals too narrow
- Alternatives?
  - Event Analysis, Interrupted Time Series: no CIs
  - ARIMA models: Differencing destroys discontinuity
  - Gaussian Process regression: CIs still too narrow
- Solution: tolerance intervals (empirical coverage), which we fit with quantile regression. Captures irreducible variance of time series, gives wide enough intervals

Findings

Netflix average daily rating: +.12 (+3%)

Netflix average daily rating: +.12 (+3%)

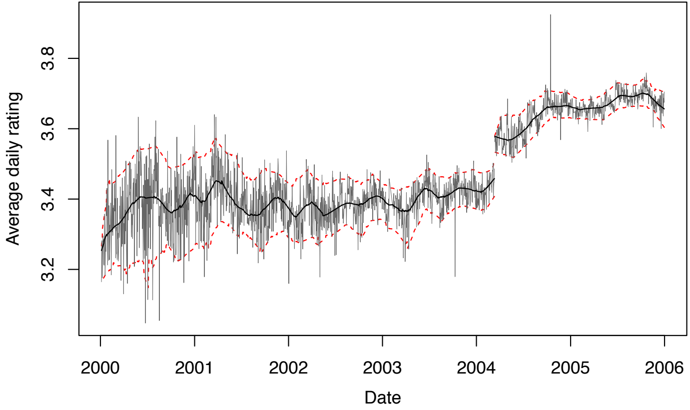

Carnegie Mellon University School of Computer Science

Facebook links: +300 new edges per day (x2)

Carnegie Mellon University School of Computer Science

Facebook links: +300 new edges per day (x2)

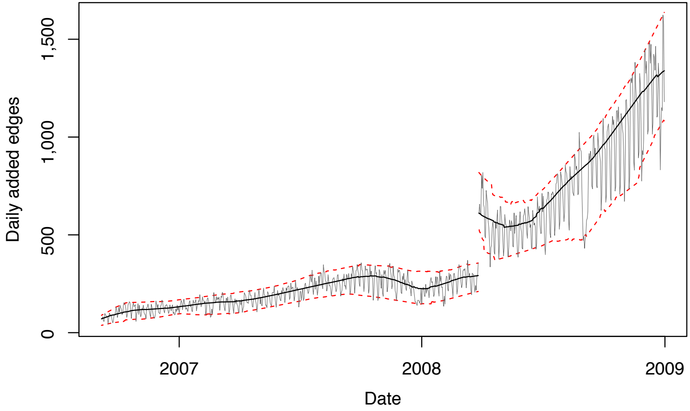

(Other Facebook New Orleans time series)

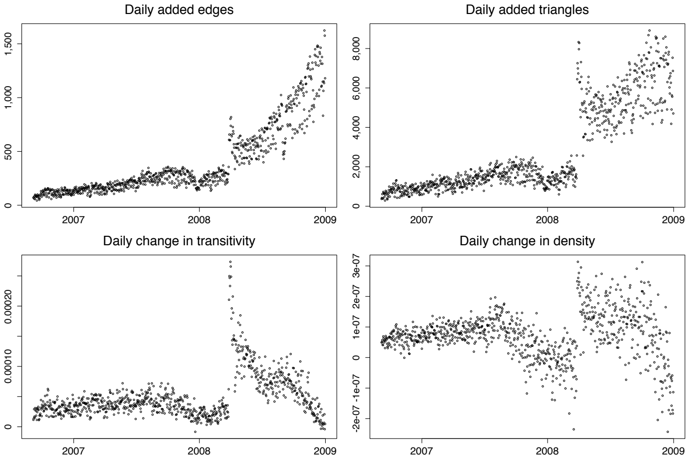

Carnegie Mellon University School of Computer Science

Facebook triangles per edge: each edge adds, on average, 3.8 more triangles (+64%)

Facebook triangles per edge: each edge adds, on average, 3.8 more triangles (+64%)

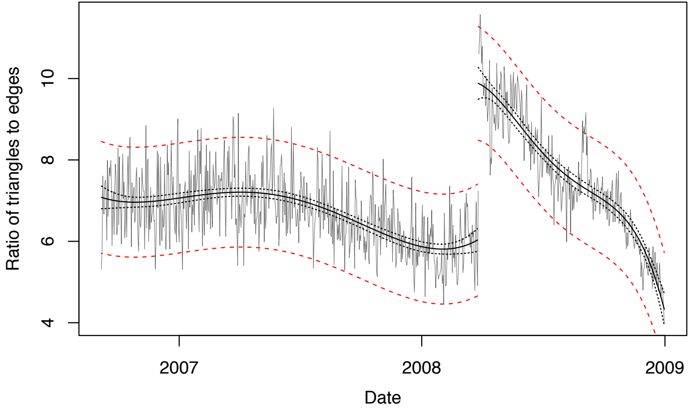

Take-aways

- Changes in the platform design can cause a massive change in user behavior—indeed, this is the intended effect of such changes!
- Specific findings:
  - What is potentially a change in response item wording (Netflix) changes user responses by about 3%
  - The addition of a triadic closure-based recommendation system doubles rate of link formation, and (initially) causes the triangle density to almost double
- "Data artifacts" are not curiosities or annoyances, they can be used to identify platform effects
- As external researchers, apply observational inference techniques to data artifacts to measure platform effects

Take-aways

- Changes in the platform design can cause a massive change in user behavior—indeed, this is the intended effect of such changes!
- Specific findings:
  - What is potentially a change in response item wording (Netflix) changes user responses by about 3%
  - The addition of a triadic closure-based recommendation system doubles rate of link formation, and (initially) causes the triangle density to almost double
- "Data artifacts" are not curiosities or annoyances, they can be used to identify platform effects
- As external researchers, apply observational inference techniques to data artifacts to measure platform effects

Take-aways

- Changes in the platform design can cause a massive change in user behavior—indeed, this is the intended effect of such changes!
- Specific findings:
  - What is potentially a change in response item wording (Netflix) changes user responses by about 3%
  - The addition of a triadic closure-based recommendation system doubles rate of link formation, and (initially) causes the triangle density to almost double
- "Data artifacts" are not curiosities or annoyances, they can be used to identify platform effects
- As external researchers, apply observational inference techniques to data artifacts to measure platform effects

Take-aways

- Changes in the platform design can cause a massive change in user behavior—indeed, this is the intended effect of such changes!
- Specific findings:
  - What is potentially a change in response item wording (Netflix) changes user responses by about 3%
  - The addition of a triadic closure-based recommendation system doubles rate of link formation, and (initially) causes the triangle density to almost double
- "Data artifacts" are not curiosities or annoyances, they can be used to identify platform effects
- As external researchers, apply observational inference techniques to data artifacts to measure platform effects

Take-aways

- Changes in the platform design can cause a massive change in user behavior—indeed, this is the intended effect of such changes!
- Specific findings:
  - What is potentially a change in response item wording (Netflix) changes user responses by about 3%
  - The addition of a triadic closure-based recommendation system doubles rate of link formation, and (initially) causes the triangle density to almost double
- "Data artifacts" are not curiosities or annoyances, they can be used to identify platform effects
- As external researchers, apply observational inference techniques to data artifacts to measure platform effects

Take-aways

- Changes in the platform design can cause a massive change in user behavior—indeed, this is the intended effect of such changes!
- Specific findings:
  - What is potentially a change in response item wording (Netflix) changes user responses by about 3%
  - The addition of a triadic closure-based recommendation system doubles rate of link formation, and (initially) causes the triangle density to almost double
- "Data artifacts" are not curiosities or annoyances, they can be used to identify platform effects
- As external researchers, apply observational inference techniques to data artifacts to measure platform effects

Take-aways

- Changes in the platform design can cause a massive change in user behavior—indeed, this is the intended effect of such changes!
- Specific findings:
  - What is potentially a change in response item wording (Netflix) changes user responses by about 3%
  - The addition of a triadic closure-based recommendation system doubles rate of link formation, and (initially) causes the triangle density to almost double
- "Data artifacts" are not curiosities or annoyances, they can be used to identify platform effects
- As external researchers, apply observational inference techniques to data artifacts to measure platform effects

Postscript

"Disciplines are revolutionized by the development of novel tools: the telescope for astronomers, **the microscope for biologists**, the particle accelerator for physicists, and brain imaging for cognitive psychologists. **Social media provide a high-powered lens into the details of human behavior and social interaction** that may prove to be equally transformative." (Golder and Macy 2012)

Carnegie Mellon University School of Computer Science

Motivation

Data

Method

Findings

Conclusions

Cells seen in 1665; cell theory in 1830s

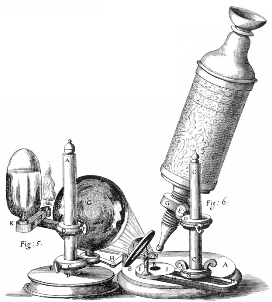

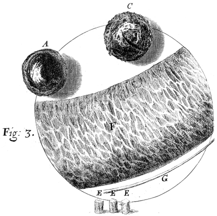

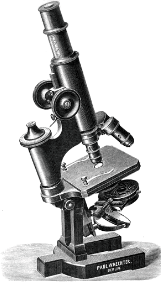

Fig. 1. Stativ I.

Paul Waechter, Optische Werkstätte, Berlin.

Fig. 4.

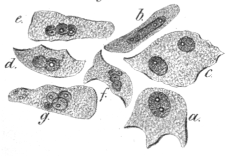

Postscript

- Don't mistake what the instrument measures for the underlying phenomenon!
- In order to know what/how to measure, we already have to have a pretty good idea of what we are looking for
- Much of science involves improving the tools as we learn more about what we are trying to study
- Social media isn't yet as high-powered a lens as we would like, but we can make it one

Postscript

- Don't mistake what the instrument measures for the underlying phenomenon!
- In order to know what/how to measure, we already have to have a pretty good idea of what we are looking for
- Much of science involves improving the tools as we learn more about what we are trying to study
- Social media isn't yet as high-powered a lens as we would like, but we can make it one

Postscript

- Don't mistake what the instrument measures for the underlying phenomenon!
- In order to know what/how to measure, we already have to have a pretty good idea of what we are looking for
- Much of science involves improving the tools as we learn more about what we are trying to study
- Social media isn't yet as high-powered a lens as we would like, but we can make it one

Postscript

- Don't mistake what the instrument measures for the underlying phenomenon!
- In order to know what/how to measure, we already have to have a pretty good idea of what we are looking for
- Much of science involves improving the tools as we learn more about what we are trying to study
- Social media isn't yet as high-powered a lens as we would like, but we can make it one

Postscript

- Don't mistake what the instrument measures for the underlying phenomenon!
- In order to know what/how to measure, we already have to have a pretty good idea of what we are looking for
- Much of science involves improving the tools as we learn more about what we are trying to study
- Social media isn't yet as high-powered a lens as we would like, but we can make it one

Carnegie Mellon University School of Computer Science

Motivation

Data

Method

Findings

Conclusions

Thank you!

momin.malik@cs.cmu.edu

mominmalik.com/icwsm2016slides.pdf