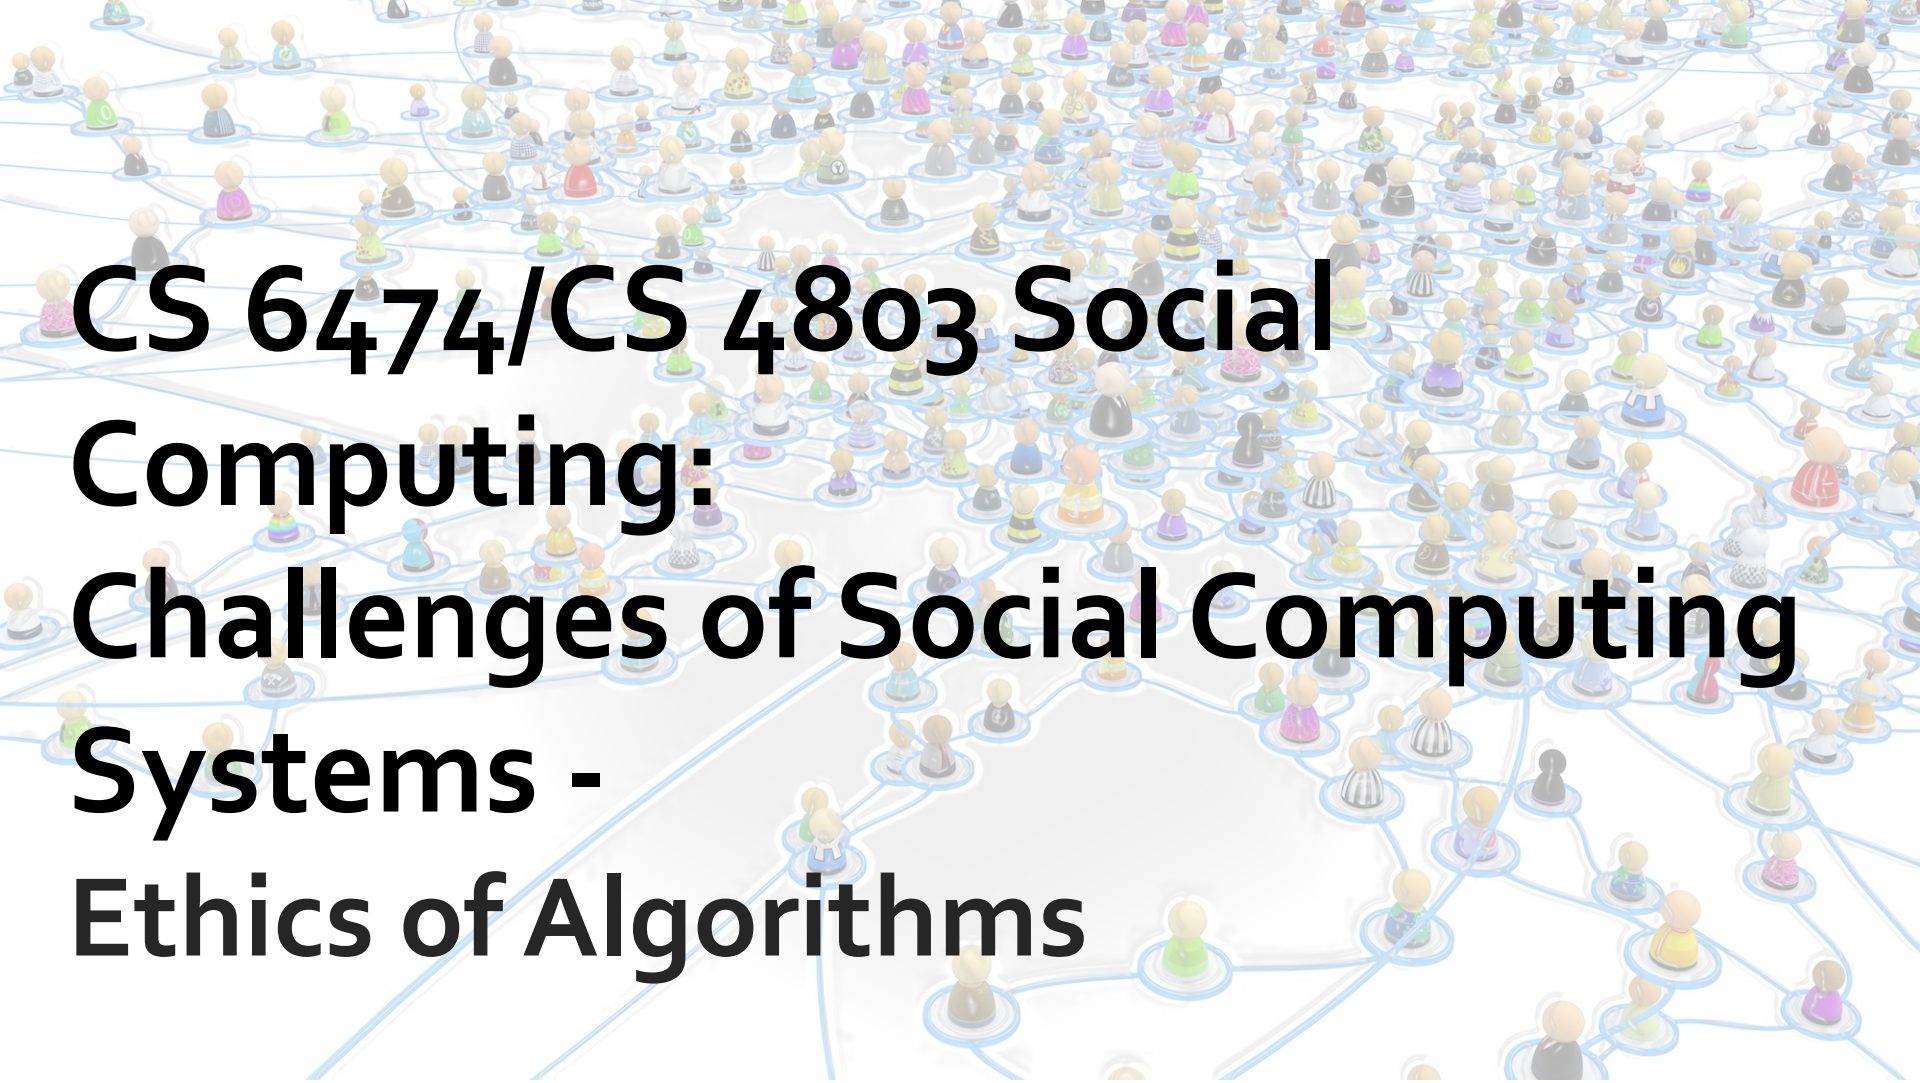

CS 6474/CS 4803 Social Computing: Challenges of Social Computing Systems - Ethics of Algorithms

Munmun De Choudhury

munmund@gatech.edu

Week 14 | April 13, 2022

Final Presentations of Term Projects

- Scheduled for Apr 20 and Apr 25
  - Signup document available on course website
- Each team gets 10 minutes in all
  - 7-8 minutes of presentation
  - 2-3 minutes of Q&A
- Each team member needs to be present
- Structure:
  - Introduction (main idea)
  - Background/Motivation/Significance
  - Related Work
  - Research questions/Goals/Objectives
  - Data/Social media platform
  - Approach/Method
  - Results
  - Implications/What you have learned/Conclusion

11 March 2013 Last updated at 15:00 ET

Share

Facebook 'likes' predict personality

Sexuality, political leanings and even intelligence can be gleaned from the things you choose to "like" on Facebook, a study suggests.

What do your Facebook likes say about you?

Researchers at Cambridge University used algorithms to predict religion, politics, race and sexual orientation.

The research, **published in the journal PNAS**, forms surprisingly accurate personal portraits, researchers said.

The findings should "ring alarm bells" for users, privacy campaigners said.

The study used 58,000 volunteers who alongside their Facebook "likes" and demographic information also provided psychometric testing results - designed to highlight personality traits.

The Facebook likes were fed into algorithms and matched with the

Related Stories

Facebook changes privacy settings

Privacy groups lobby Facebook

Top stories

Soldier killed in Canada shootings

WHO crisis team holds Ebola talks

Sweden could use force in sub hunt

U-boat found with ship it sank

Total plane crash driver 'got lost'

"This research should ring alarm bells for anyone who thinks that privacy settings are the solution to protecting information online. We need to fundamentally re-think how much data we are voluntarily sharing," said Nick Pickles, director of privacy campaign group Big Brother Watch.

"Yet again, it is clear the lack of transparency about how users' data is being used will lead to entirely justified fears about our data being exploited for commercial gain."

Deep neural networks are more accurate than humans at detecting sexual orientation from facial images.

Yilun Wang, Michal Kosinski

Created on: September 07, 2017 | Last edited: October 16, 2017

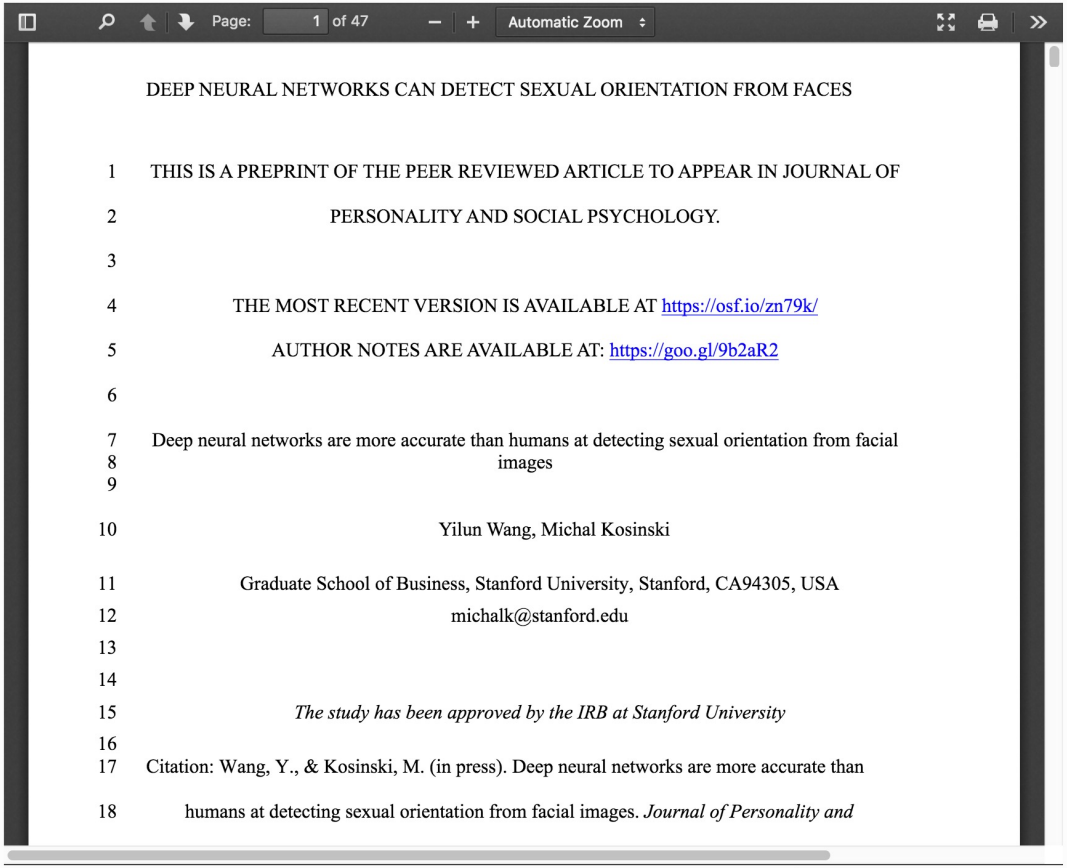

DEEP NEURAL NETWORKS CAN DETECT SEXUAL ORIENTATION FROM FACES

THIS IS A PREPRINT OF THE PEER REVIEWED ARTICLE TO APPEAR IN JOURNAL OF PERSONALITY AND SOCIAL PSYCHOLOGY.

THE MOST RECENT VERSION IS AVAILABLE AT https://osf.io/zn79k/

AUTHOR NOTES ARE AVAILABLE AT: https://goo.gl/9b2aR2

Deep neural networks are more accurate than humans at detecting sexual orientation from facial images

Yilun Wang, Michal Kosinski

Graduate School of Business, Stanford University, Stanford, CA94305, USA
michalk@stanford.edu

The study has been approved by the IRB at Stanford University

Citation: Wang, Y., & Kosinski, M. (in press). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of Personality and*

wang_kosinski.pdf

Download preprint Version: 10

Download preprint Downloads: 24829

Abstract

We show that faces contain much more information about sexual orientation than can be perceived and interpreted by the human brain. We used deep neural networks to extract features from 35,326 facial images. These features were entered into a logistic regression aimed at classifying sexual orientation. Given a single facial image, a classifier ...

See more

Preprint DOI

10.17605/OSF.IO/HV28A

License

Disciplines

Social and Behavioral Sciences | Psychology

Tags

Artificial Intelligence | Big Data | Computational Social Science | Computer Vision | Facial recognition | Prenatal Hormone Theory | Privacy | Sexual orientation

Citations

APA

Wang, Y., & Kosinski, M. (2017, October 16). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images.. Retrieved from psyarxiv.com/hv28a

Automatic Crime Prediction using Events Extracted from Twitter Posts

Xiaofeng Wang, Matthew S. Gerber, and Donald E. Brown

Department of Systems and Information Engineering, University of Virginia
{xw4u,msg8u,brown}@virginia.edu

Abstract. Prior work on criminal incident prediction has relied primarily on the historical crime record and various geospatial and demographic information sources. Although promising, these models do not take into account the rich and rapidly expanding social media context that surrounds incidents of interest. This paper presents a preliminary investigation of Twitter-based criminal incident prediction. Our approach is based on the automatic semantic analysis and understanding of natural language Twitter posts, combined with dimensionality reduction via latent Dirichlet allocation and prediction via linear modeling. We tested our model on the task of predicting future hit-and-run crimes. Evaluation results indicate that the model comfortably outperforms a baseline model that predicts hit-and-run incidents uniformly across all days.

1 Introduction

Traditional crime prediction systems (e.g., the one described by Wang and Brown [14]) make extensive use of historical incident patterns as well as layers of in-

1. What can we do with data generated from social computing systems? What can we study?

2. What should we **not** do with these data. What study designs are particularly **problematic**?

Challenges

Experimental evidence of massive-scale emotional contagion through social networks

Summary

- This experiment was widely criticized on ethical grounds regarding informed consent.

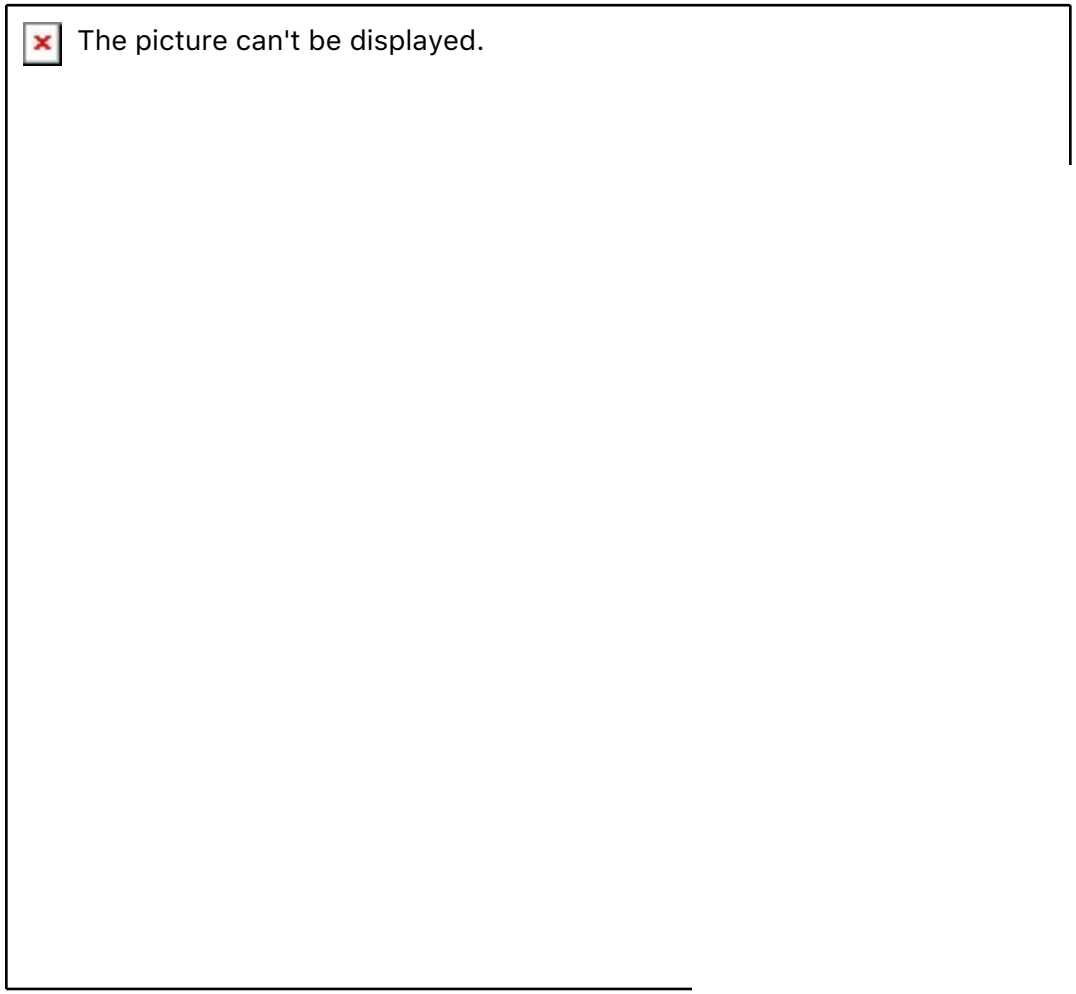

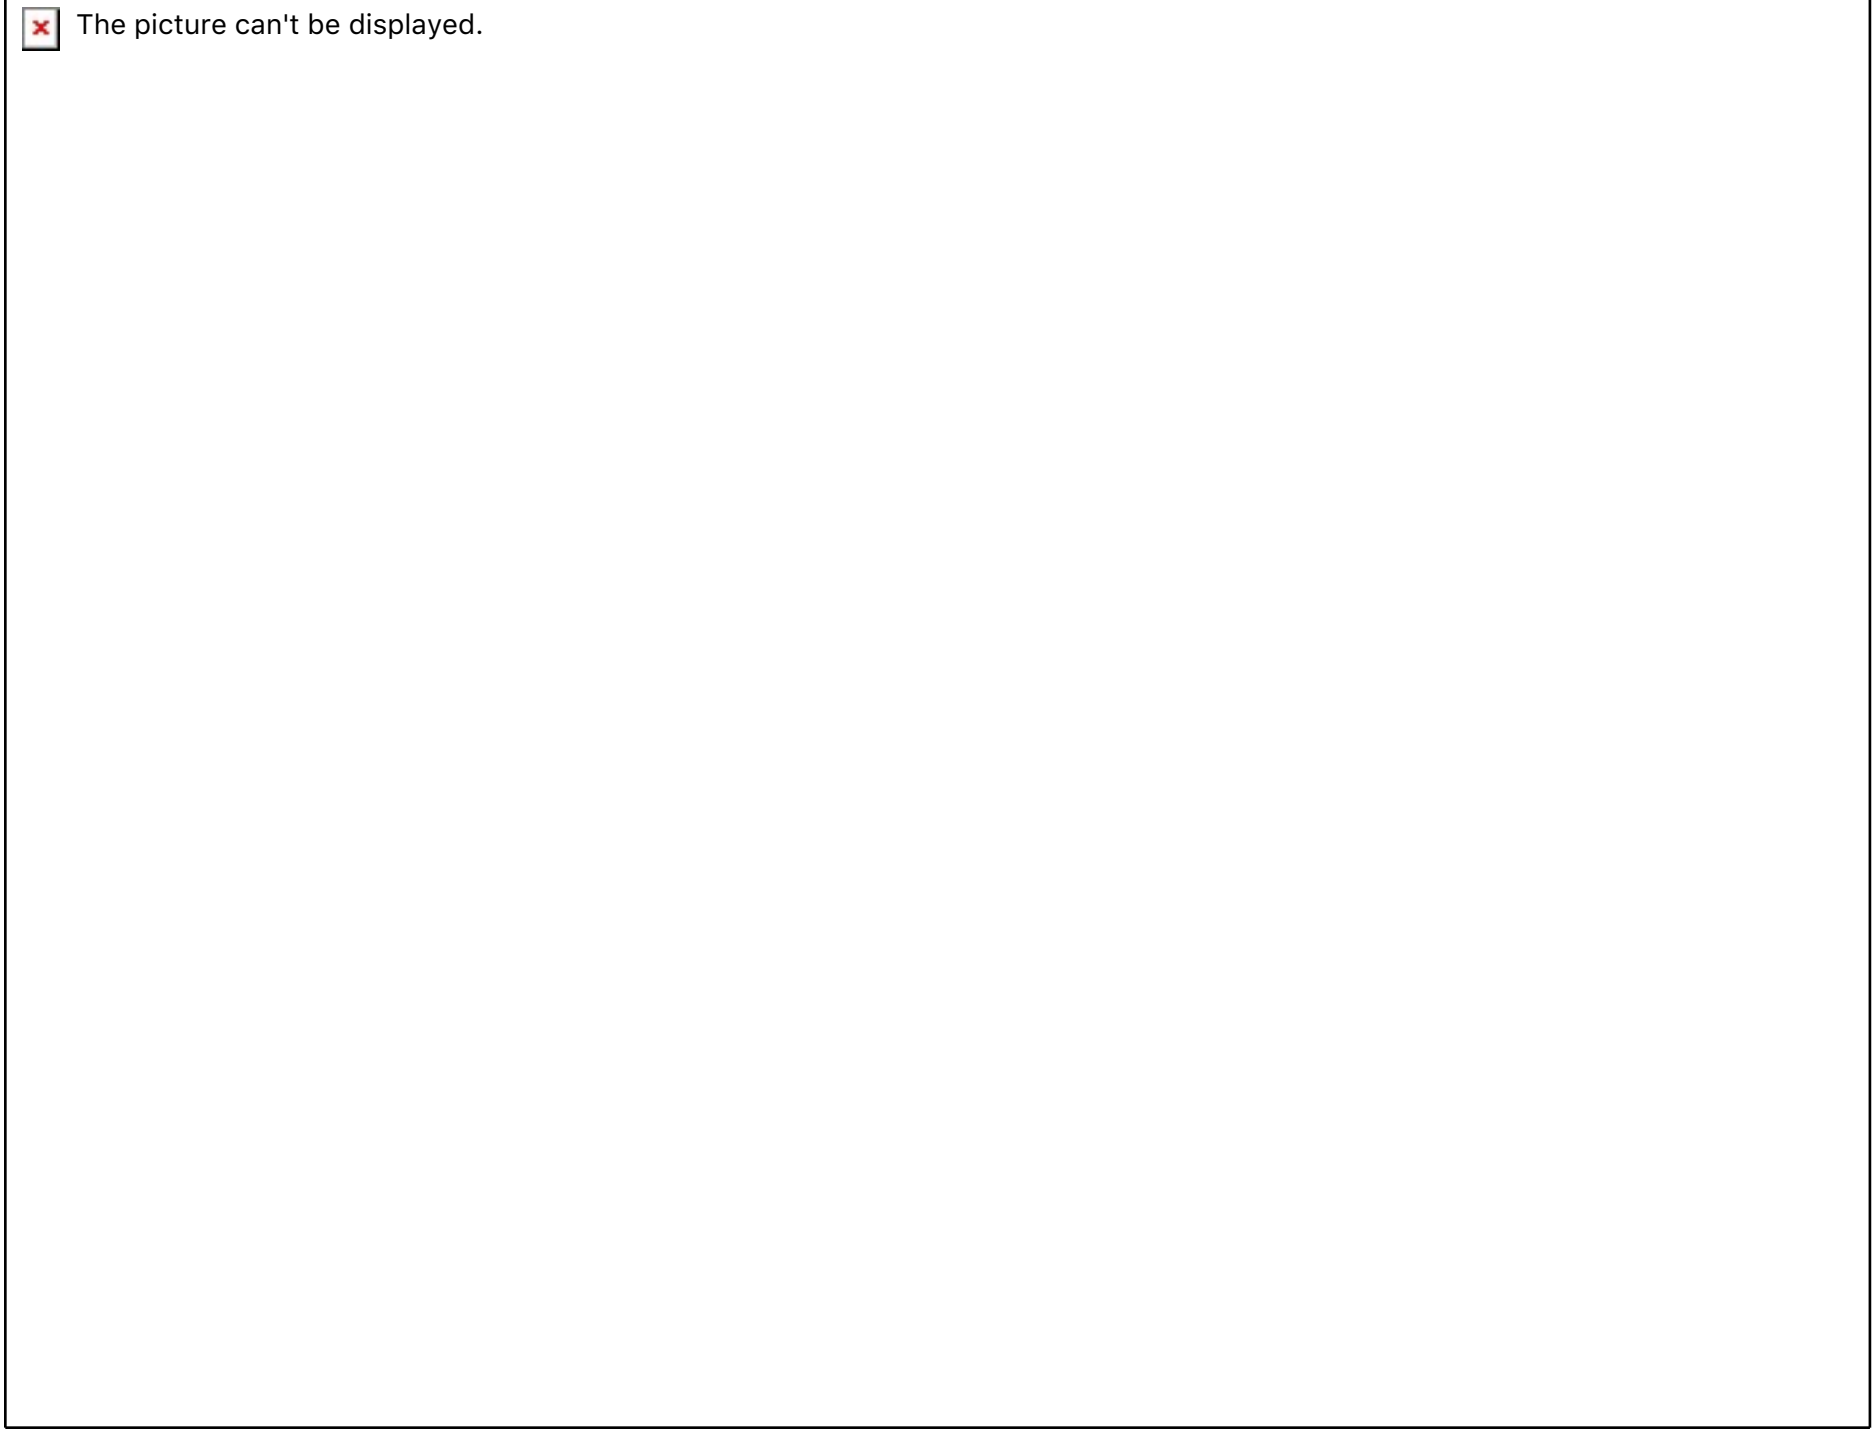
The picture can't be displayed.

Highlights of some findings…

- **Living in a lab**
  - *Dear Mr. Zuckerburg, Last I checked, we did not decide to jump in a petri dish to be utilized at your disposal . . . We connect with our loved ones.*
- ***Manipulation anxieties***
  - *Don't be fooled, manipulating a mood is the ability to manipulate a mind. Political outcomes, commerce, and civil unrest are just a short list of things that can be controlled.*
- **Wake up, sheeple**
  - *Anyone who doesn't realise that anything you put "out there" on Facebook (or any other social media site) is like shouting it through a bullhorn should have their internet competency licence revoked. We can't blame all stupidity on some or other conspiracy…*
- **No big deal**
  - *A/B testing (i.e. basically what happened here) when software companies change content or algorithms for a subset of users happens *all the time*. It's standard industry practice.*

A key takeaway –
consent is important!

Consent at Scale – why it is hard


The picture can't be displayed.

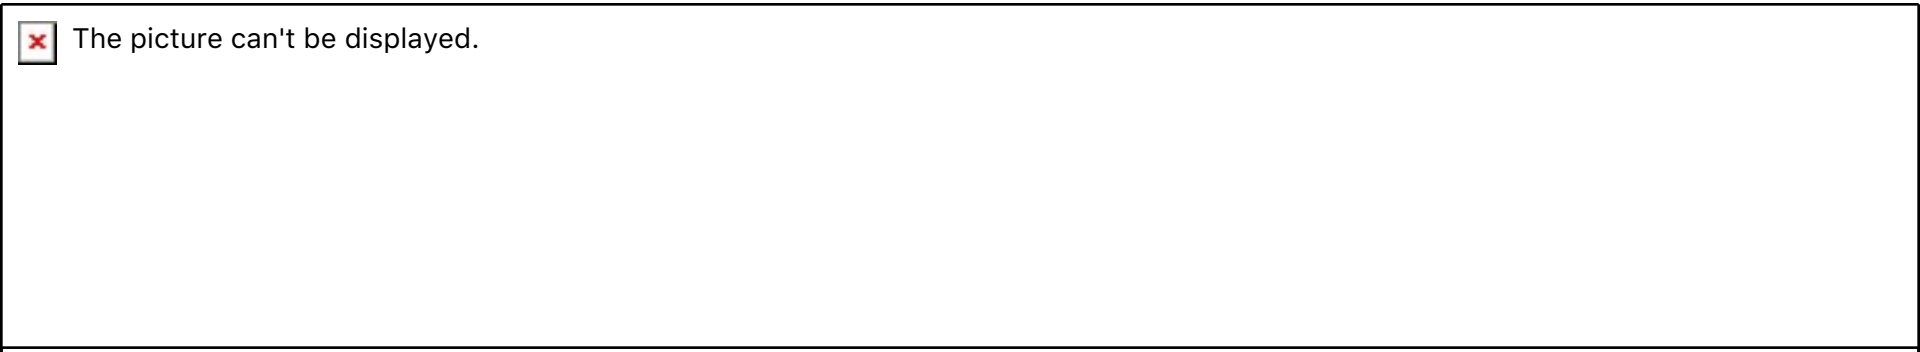
The picture can't be displayed.

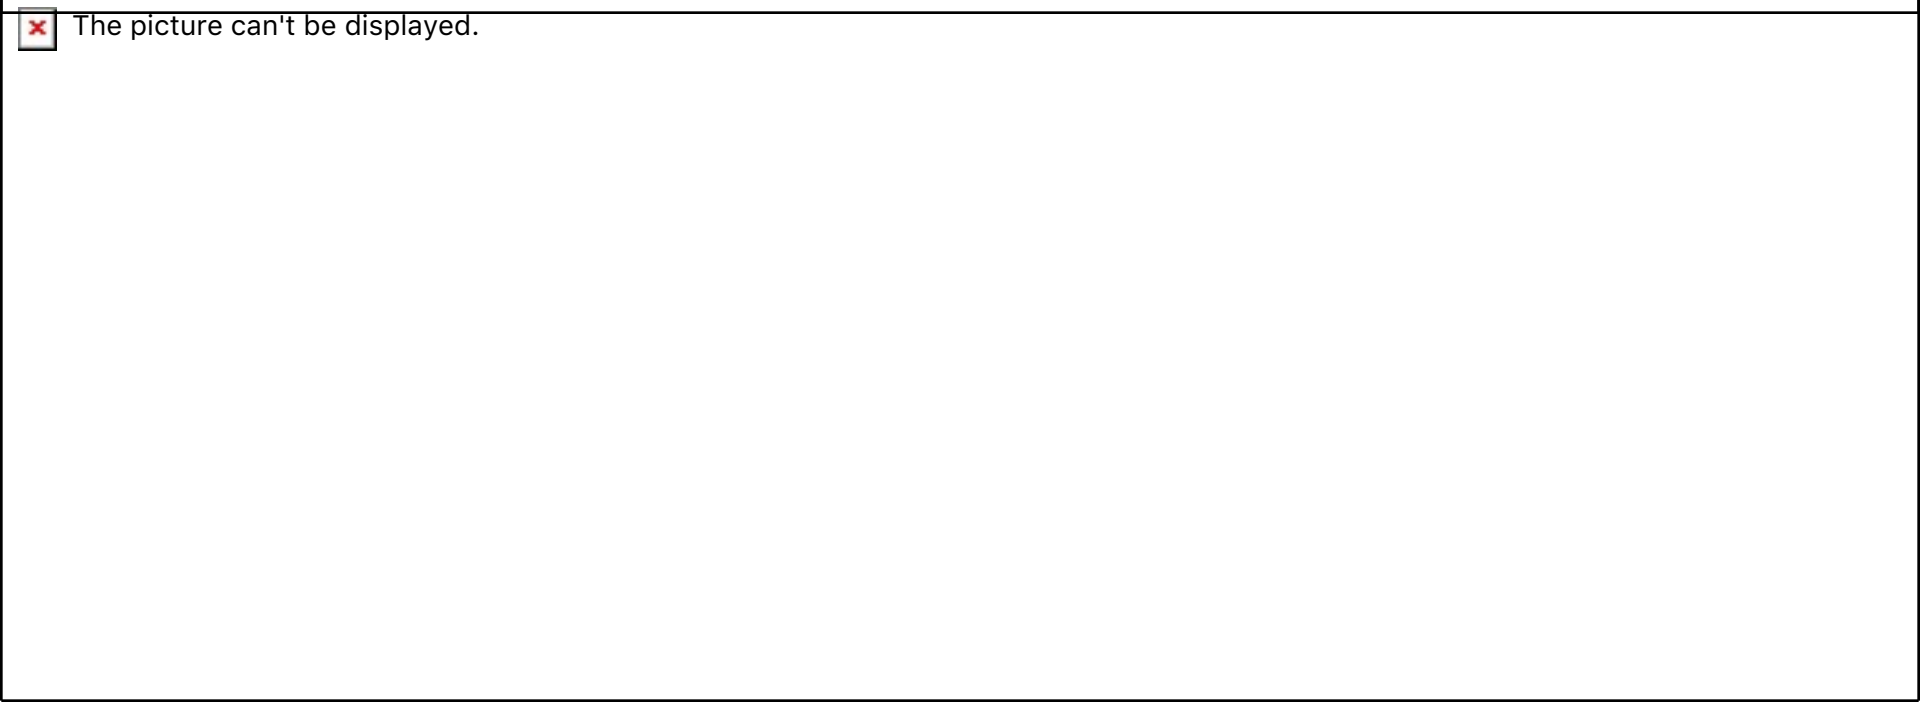
The picture can't be displayed.

The Case of Deleted Tweets/Social media posts

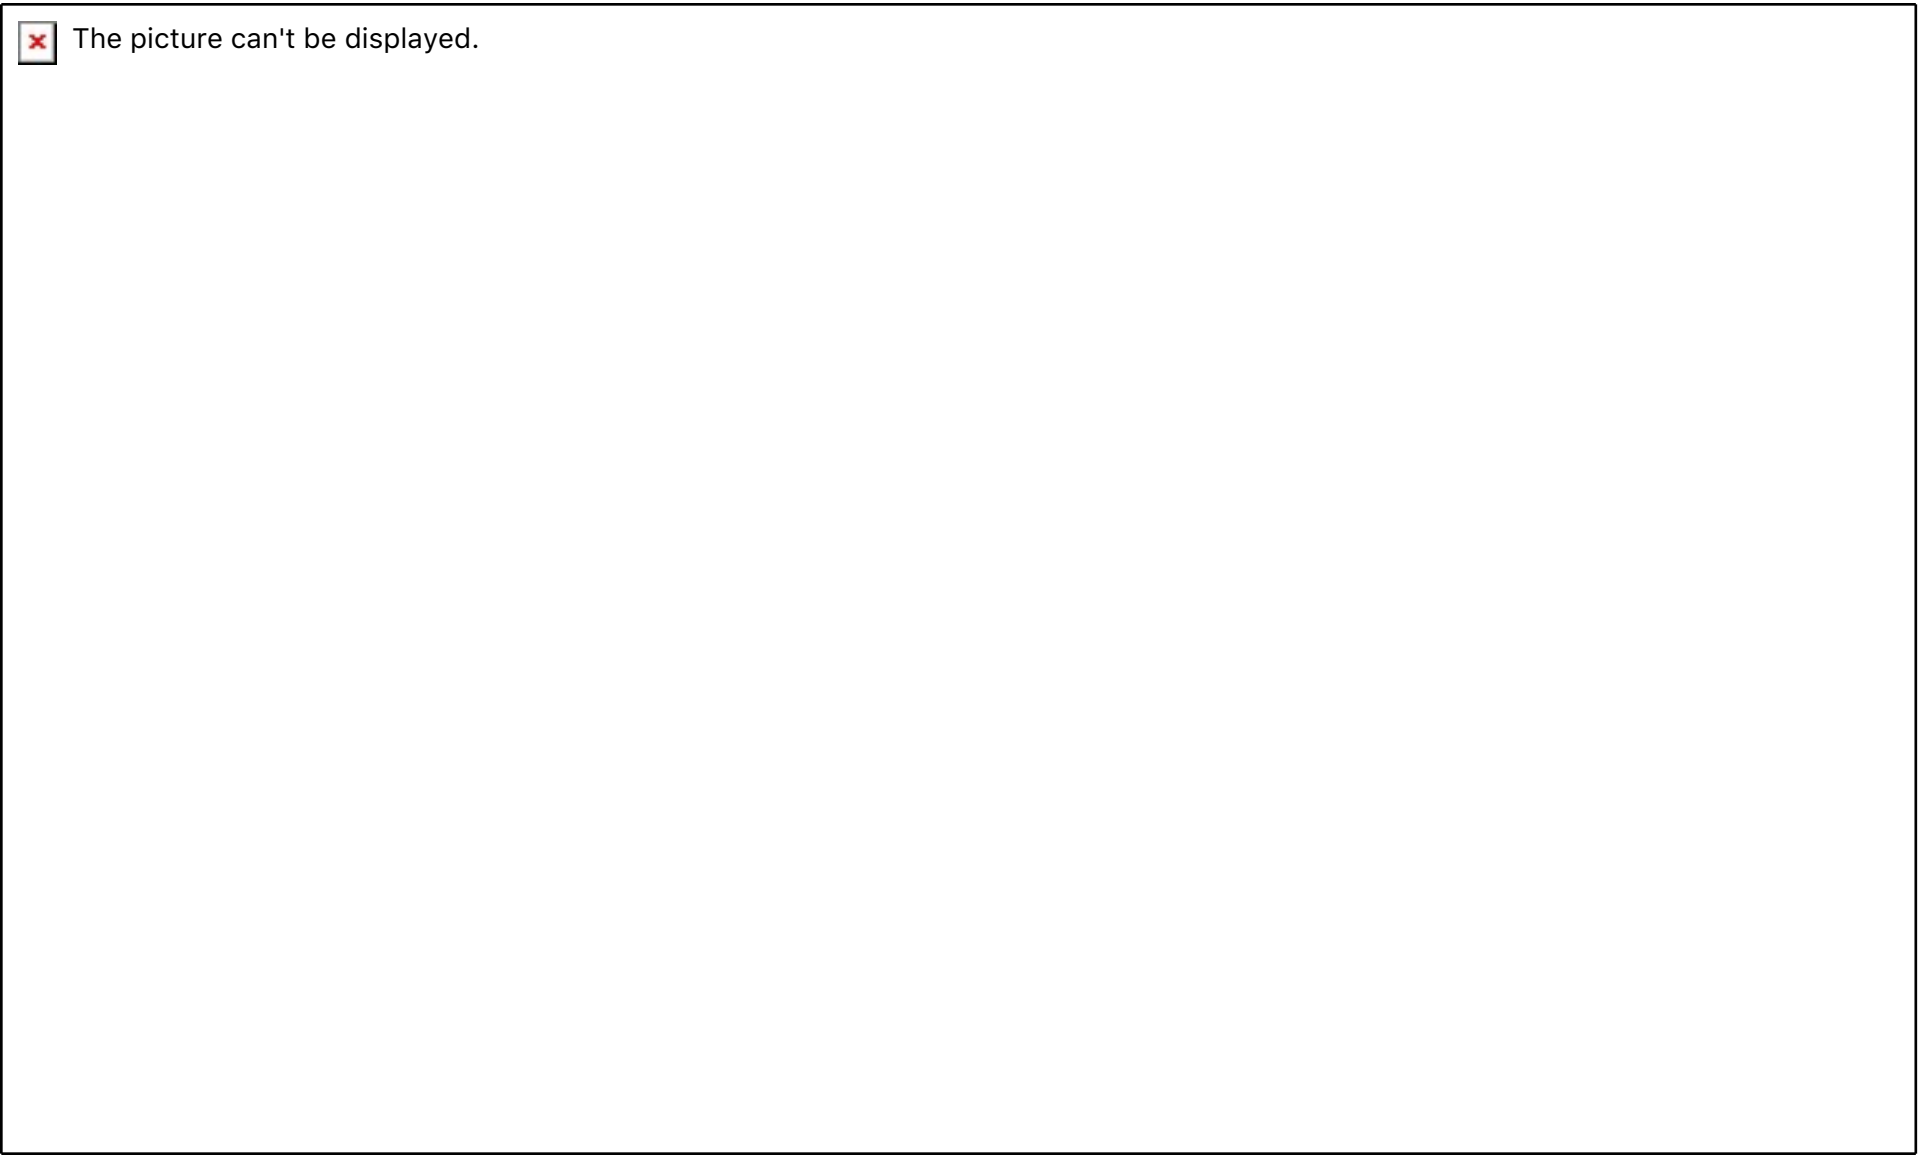

Also what about those who can‘t give consent any more? *The case of dead people*

- Today‘s view
- Medieval view

- Things are muddled when it comes to dead people‘s digital lives – legislation has not kept up with technological change

Digital Wills and Beneficiaries (Forbes)

… still particularly nascent when it comes to data stored by a third-party company

https://www.forbes.com/sites/bernardmarr/2017/02/01/what-really-happens-to-your-big-data-after-you-die/

When there is no consent, researchers have poor understanding of what can go wrong, and “participants” or research subjects have limited understanding of risk.


The picture can't be displayed.

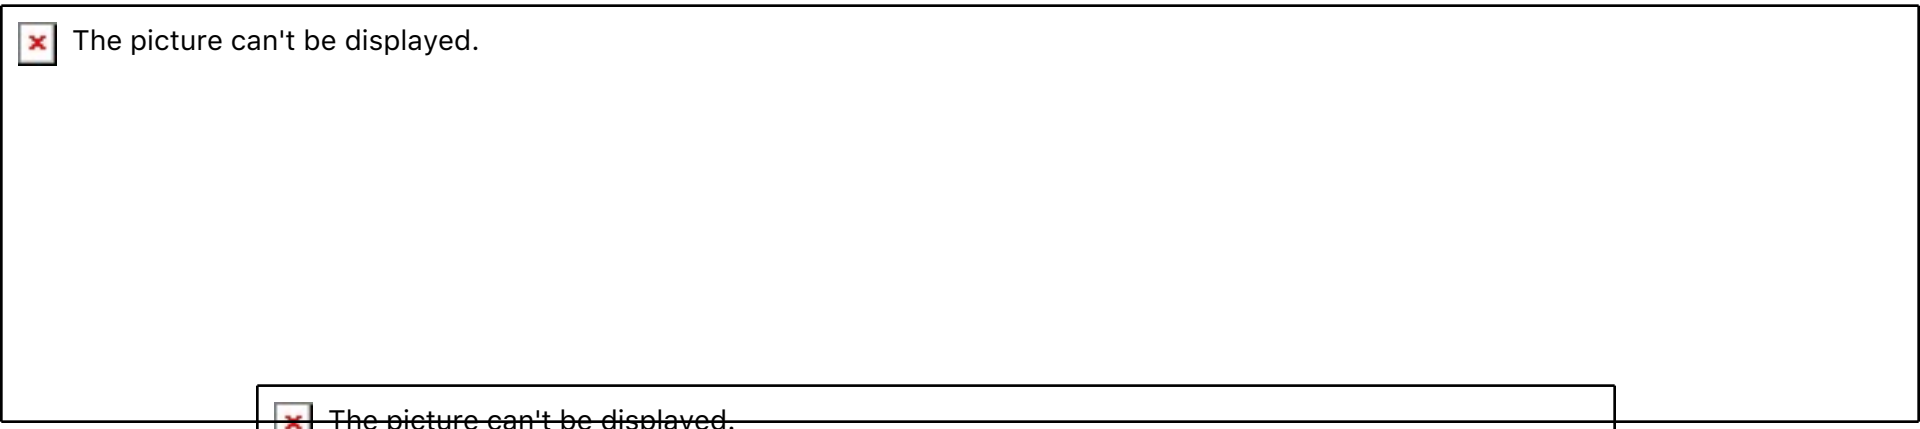
The picture can't be displayed.

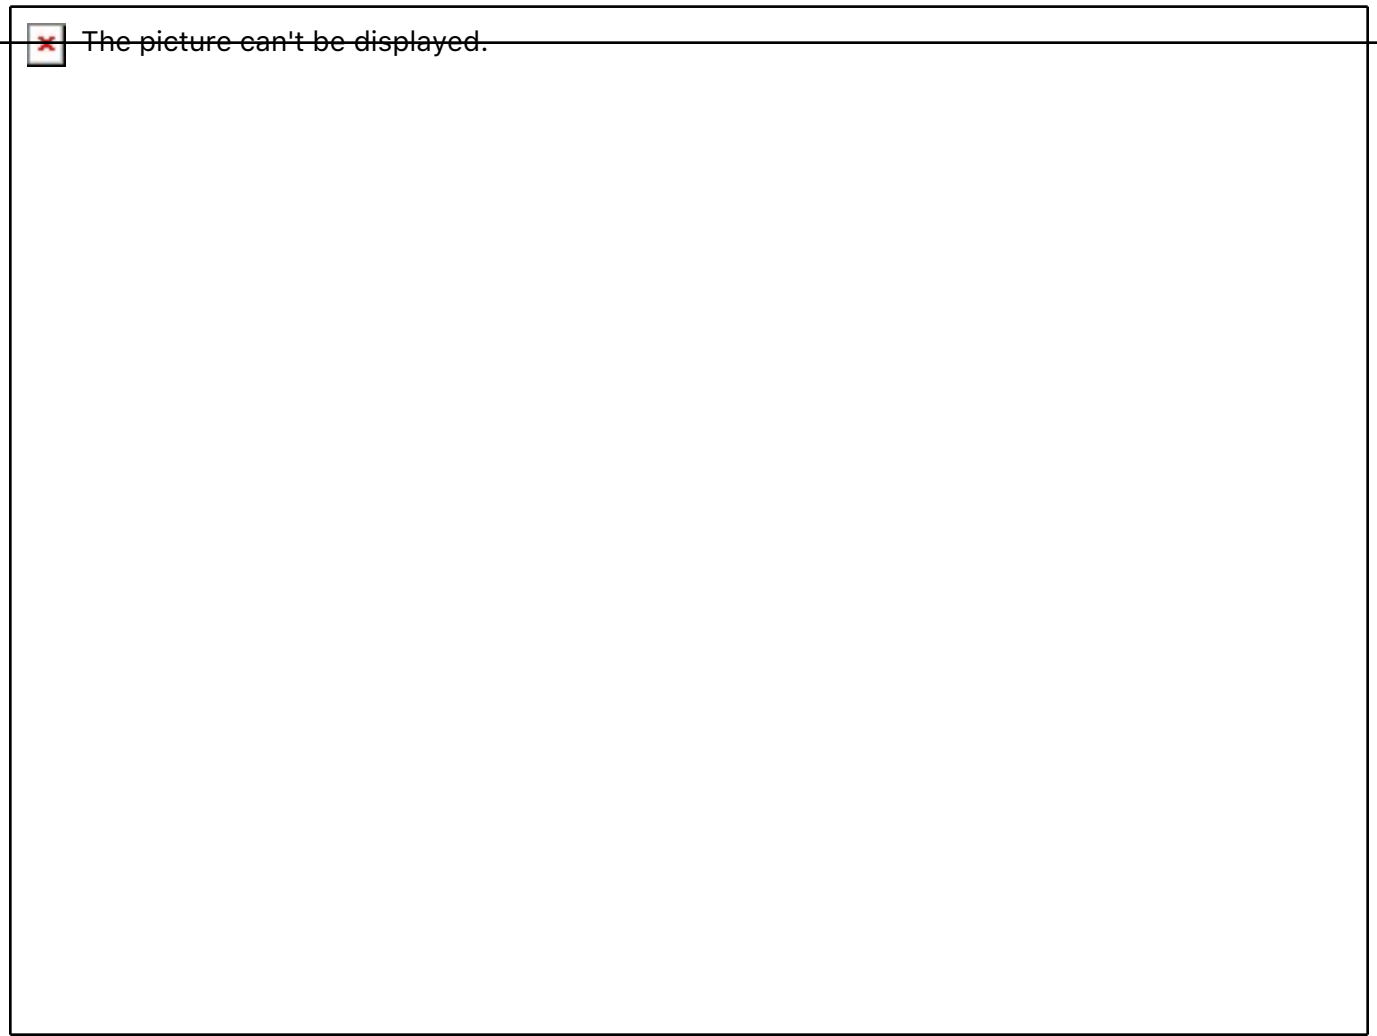
The picture can't be displayed.


The picture can't be displayed.

Amid the ethical controversy surrounding the experiment, Facebook twice attempted to draw attention to the study's claims about well-being. Lead author Adam Kramer wrote:

> *The reason we did this research is because we care about the emotional impact of Facebook and the people that use our product. We felt that it was important to investigate the common worry that seeing friends post positive content leads to people feeling negative or left out. … And we found the exact opposite to what was then the conventional wisdom: Seeing a certain kind of emotion (positive) encourages it rather than suppresses is [sic]. (2014)*

Mike Schroepfer, Facebook's Chief Technology Officer, later reiterated Kramer's statement (2014).

Class Exercise I

Redo the emotion contagion study experimentally or with observational data, but in an ethical manner. What study design will you use?

If true, these findings could substantially alleviate concern that Facebook represents a threat to well-being. But the work also has significant methodological concerns.

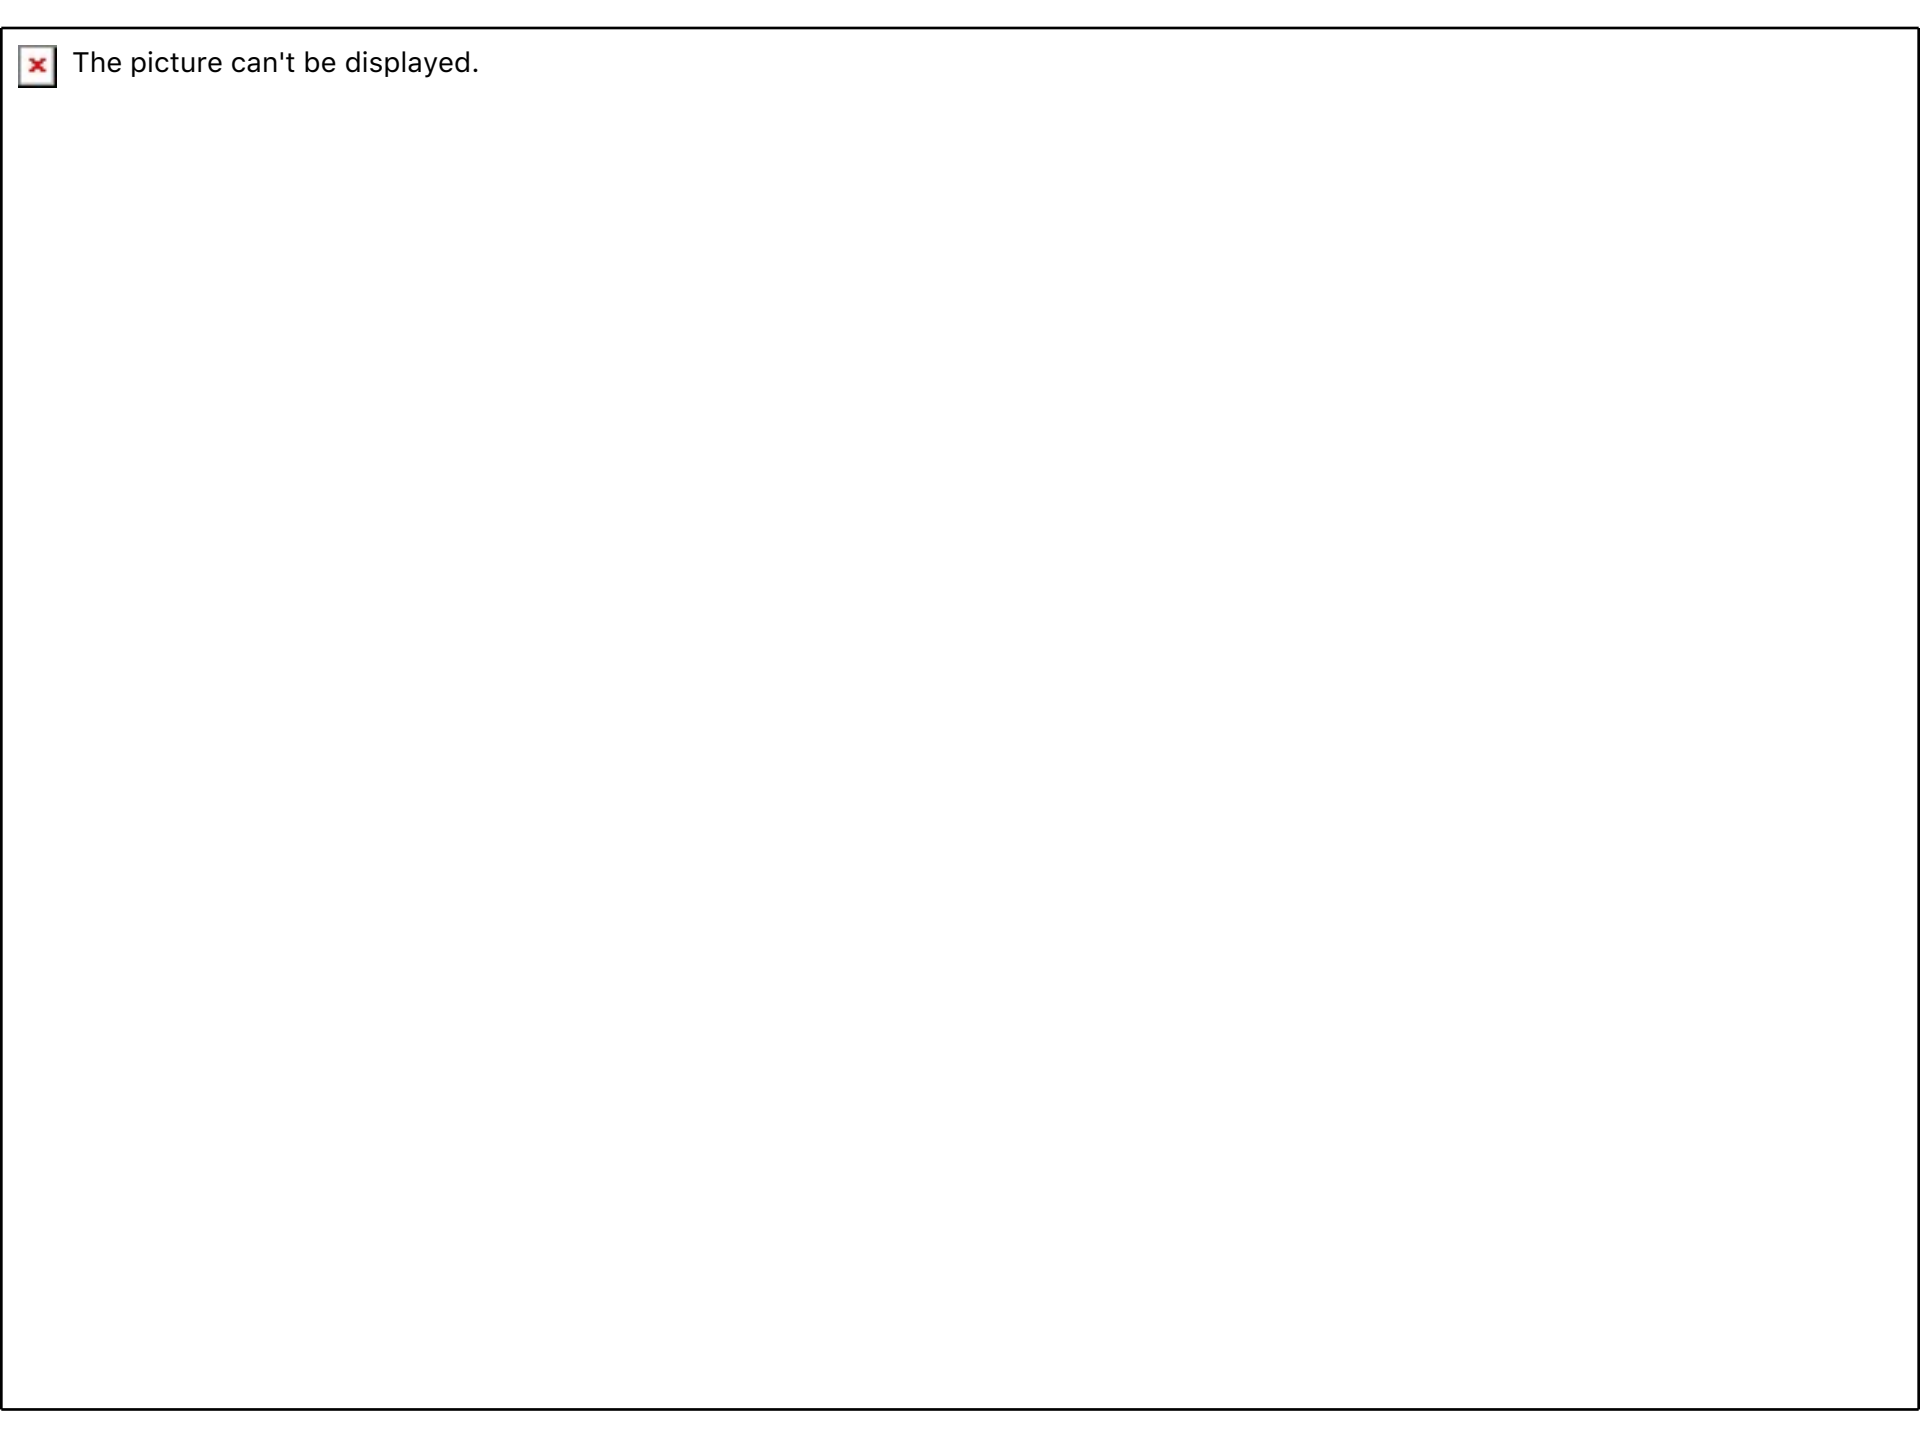
The picture can't be displayed.

The picture can't be displayed.

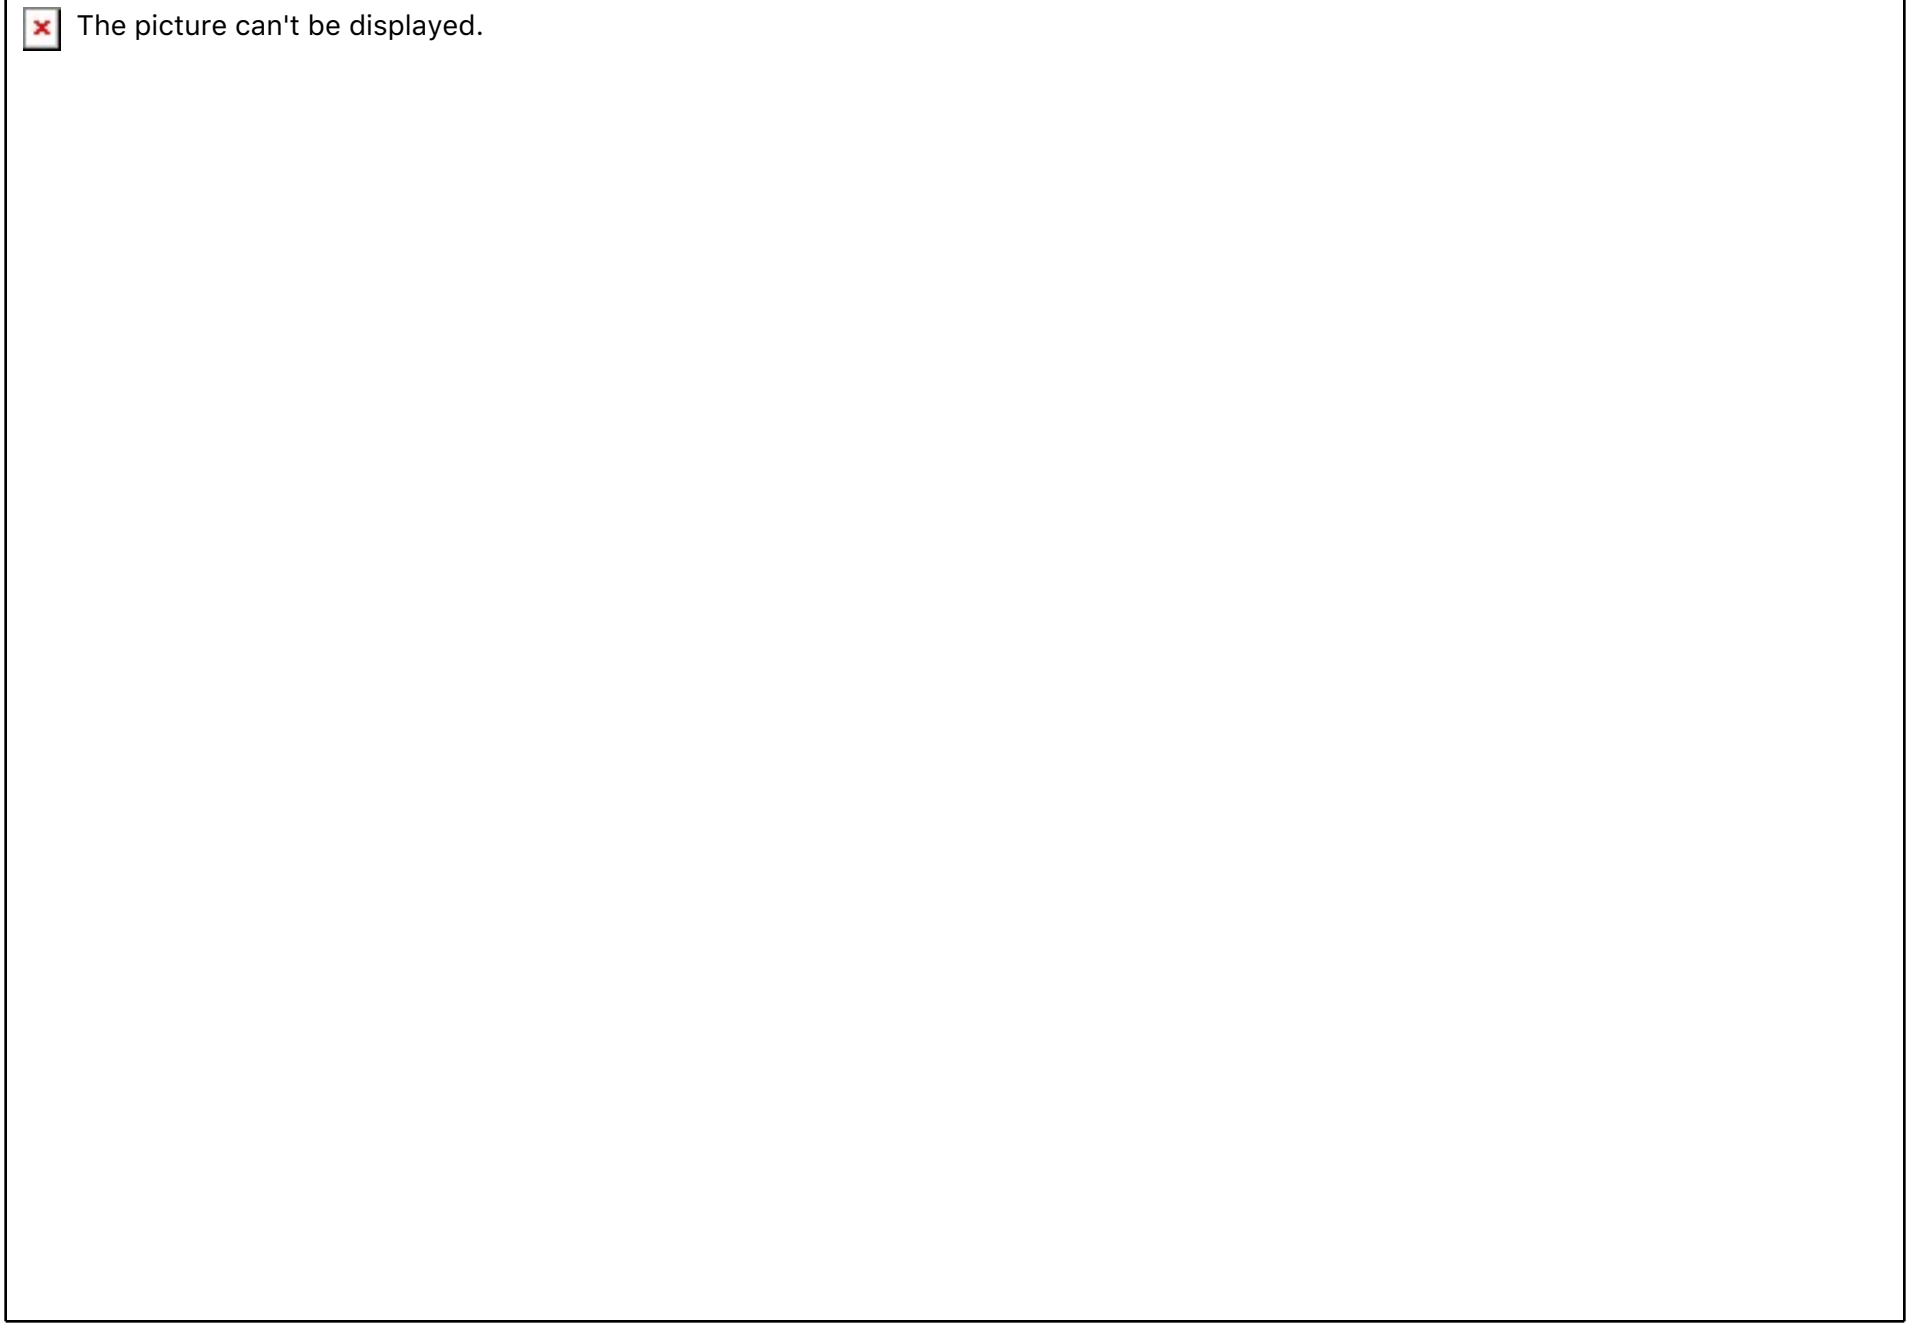
The picture can't be displayed.

The picture can't be displayed.

A Taxonomy of Ethical Tensions in Inferring Mental Health States from Social Media

Overview of Taxonomy

- Participant and research oversight
- Validity, interpretability, and methods
- Stakeholder implications

Possible Ethical Solutions

Class Exercise II

Analyze the challenges in the recently released suicide prevention AI tool of Facebook. Analyze from the perspective of 1) informed consent; 2) methodology/algorithm; and 3) transparency.