Adversarial Removal of Demographic Attributes from Text Data

Yanai Elazar and Yoav Goldberg

Bar-Ilan University / NLP Group

November 2, 2018

אוניברסיטת בר-אילן
Bar-Ilan University

BIU NLP

Motivation

Text is used for predictions

Motivation

- For example, consider a text classification setup, where we predict:
  - Hiring decisions
  - Mortgages approval
  - Loans rates

Motivation

Department of Linguistics &
Department of Computer Science,
Stanford University
Stanford CA 94305-2150

Education

B.A Linguistics, with honors, University of California at Berkeley, 1983

Ph.D. Computer Science, University of California at Berkeley, 1992

Postdoc, International Computer Science Institute, Berkeley, 1992-1995

Academic Employment

Stanford University: Professor and Chair of Linguistics and Professor of Computer Science, 2014-

Stanford University: Professor of Linguistics and (by courtesy) of Computer Science, 2010-

Stanford University: Associate Professor of Linguistics and (by courtesy) of Computer Science, 2004-2010

University of Colorado, Associate Professor of Linguistics, Computer Science, Cognitive Science, 2001-2003

University of Colorado, Assistant Professor of Linguistics, Computer Science, and Cognitive Science, 1996-2001

This applicant would easily get any NLP job

Motivation

The common implementation:

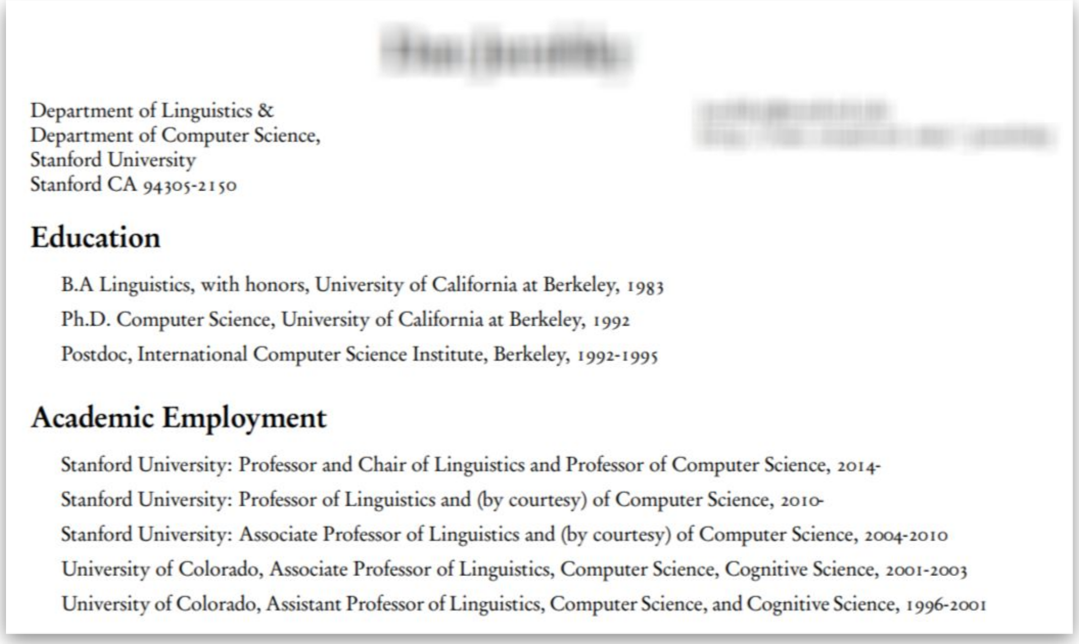
Department of Linguistics &
Department of Computer Science,
Stanford University
Stanford CA 94305-2150

Education

B.A Linguistics, with honors, University of California at Berkeley, 1983
Ph.D. Computer Science, University of California at Berkeley, 1992
Postdoc, International Computer Science Institute, Berkeley, 1992-1995

Academic Employment

Stanford University: Professor and Chair of Linguistics and Professor of Computer Science, 2014-
Stanford University: Professor of Linguistics and (by courtesy) of Computer Science, 2010-
Stanford University: Associate Professor of Linguistics and (by courtesy) of Computer Science, 2004-2010
University of Colorado, Associate Professor of Linguistics, Computer Science, Cognitive Science, 2001-2003
University of Colorado, Assistant Professor of Linguistics, Computer Science, and Cognitive Science, 1996-2001

Input CV

ML
Model

Hire

Don't Hire

Motivation

The common implementation:

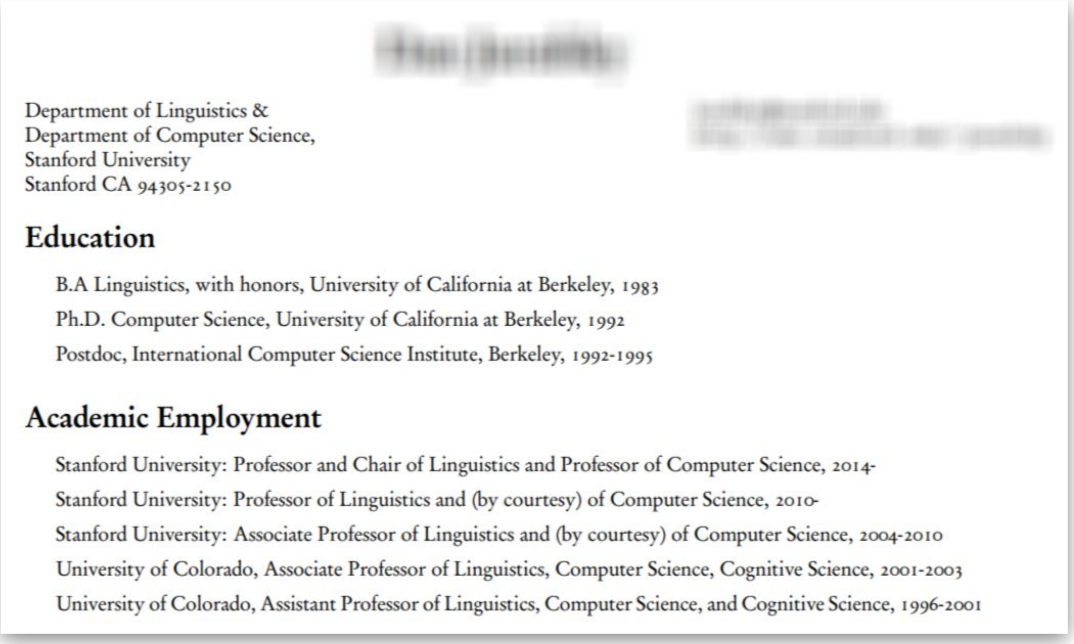

Department of Linguistics &
Department of Computer Science,
Stanford University
Stanford CA 94305-2150

Education

B.A Linguistics, with honors, University of California at Berkeley, 1983
Ph.D. Computer Science, University of California at Berkeley, 1992
Postdoc, International Computer Science Institute, Berkeley, 1992-1995

Academic Employment

Stanford University: Professor and Chair of Linguistics and Professor of Computer Science, 2014-
Stanford University: Professor of Linguistics and (by courtesy) of Computer Science, 2010-
Stanford University: Associate Professor of Linguistics and (by courtesy) of Computer Science, 2004-2010
University of Colorado, Associate Professor of Linguistics, Computer Science, Cognitive Science, 2001-2003
University of Colorado, Assistant Professor of Linguistics, Computer Science, and Cognitive Science, 1996-2001

Input CV

$\xrightarrow{\text{Encode}}$ Representation $\xrightarrow{\text{Predict}}$

Hire

Don’t Hire

Motivation

- But then we see this

BUSINESS INSIDER

TECH | FINANCE | POLITICS | STRATEGY | LIFE | ALL — PRIME | INTELLIGENCE

Amazon built an AI tool to hire people but had to shut it down because it was discriminating against women

Isobel Asher Hamilton 19h

CNBC

MENU — MARKETS — BUSINESS NEWS — INVESTING — TECH — POLITICS — CNBC TV

RETAIL

APPAREL | DISCOUNTERS | DEPARTMENT STORES | E-COMMERCE | FOOD AND BEVERA

Amazon scraps a secret A.I. recruiting tool that showed bias against women

:d a big problem: their new

e 2014 to review job
search for top talent, five

l intelligence to give job
uch like shoppers rate

ct 2018

REUTERS

BUSINESS NEWS — OCTOBER 10, 2018 / 6:12 AM / UPDATED 16 HOURS AGO

Amazon scraps secret AI recruiting tool that showed bias against women

Motivation

- When deciding on recruiting an applicant based on their writings/CV...
- ...we would like that attributes like the author’s:
  - Gender
  - Race
  - Age
- won’t be part of the decision.
- In some places, this is even illegal

Motivation

- We seek to build models which are:
  - Predictive for some main task (e.g. Hiring decision)

  - Agnostic to irrelevant/protected attributes (e.g. race, gender, …)

Motivation

How do we know we do not condition on some sensitive attribute by mistake?

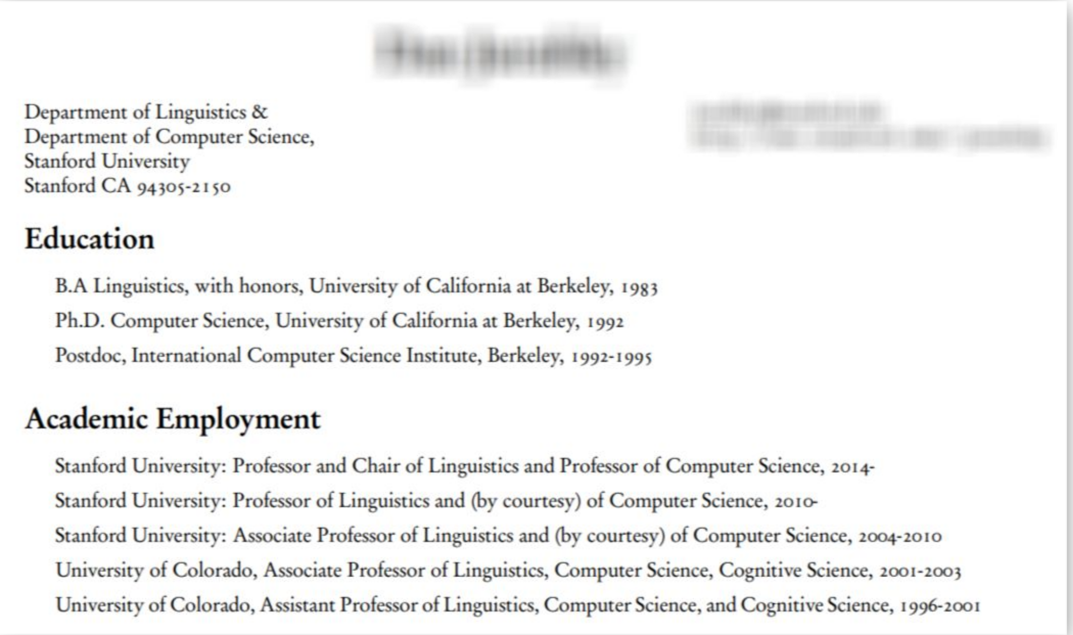
Department of Linguistics &
Department of Computer Science,
Stanford University
Stanford CA 94305-2150

Education

B.A Linguistics, with honors, University of California at Berkeley, 1983
Ph.D. Computer Science, University of California at Berkeley, 1992
Postdoc, International Computer Science Institute, Berkeley, 1992-1995

Academic Employment

Stanford University: Professor and Chair of Linguistics and Professor of Computer Science, 2014-
Stanford University: Professor of Linguistics and (by courtesy) of Computer Science, 2010-
Stanford University: Associate Professor of Linguistics and (by courtesy) of Computer Science, 2004-2010
University of Colorado, Associate Professor of Linguistics, Computer Science, Cognitive Science, 2001-2003
University of Colorado, Assistant Professor of Linguistics, Computer Science, and Cognitive Science, 1996-2001

Input CV

Encode → Representation → Predict

Hire

Don't Hire

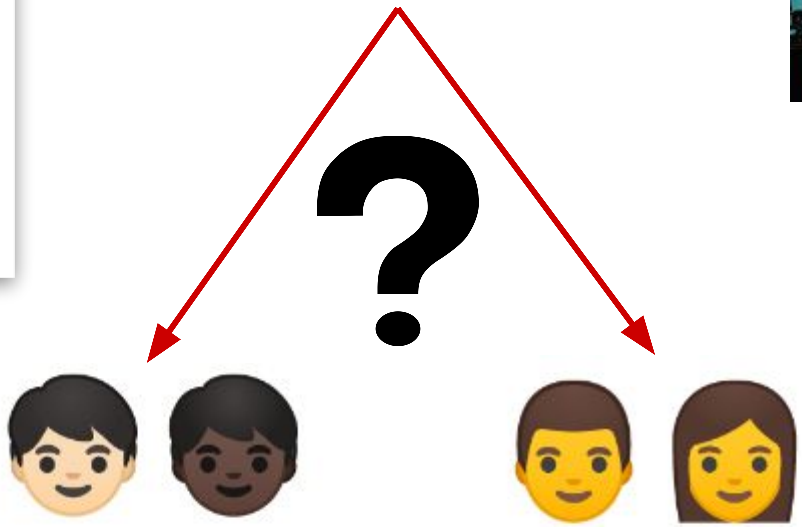

Motivation

If we **can** predict protected attributes from the representation...

A talented candidate might suffer from demographic discrimination

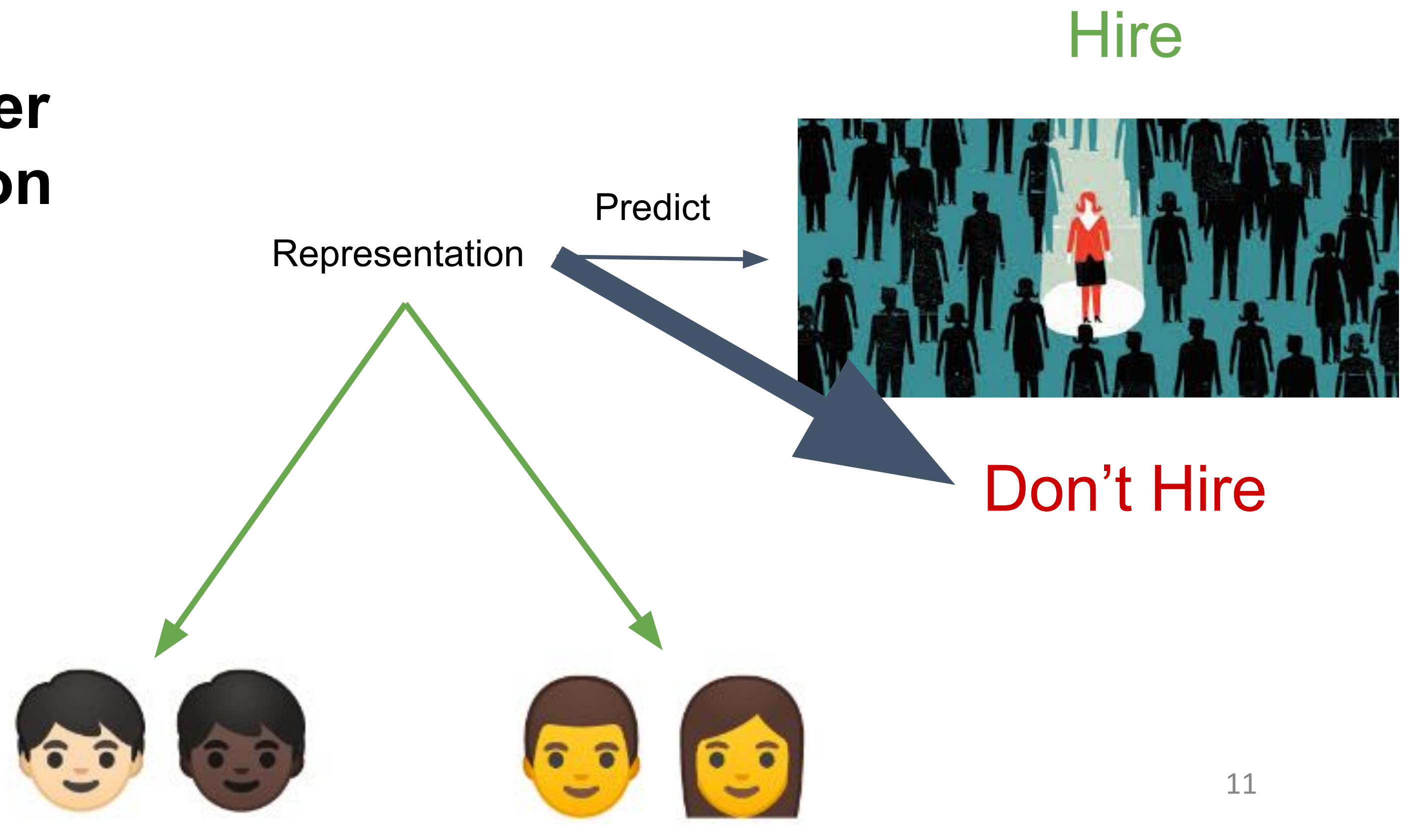

Motivation

If we **can not** predict protected attributes from the representation...

We don’t condition on these protected attributes and...
A talented candidate won’t suffer from demographic discrimination

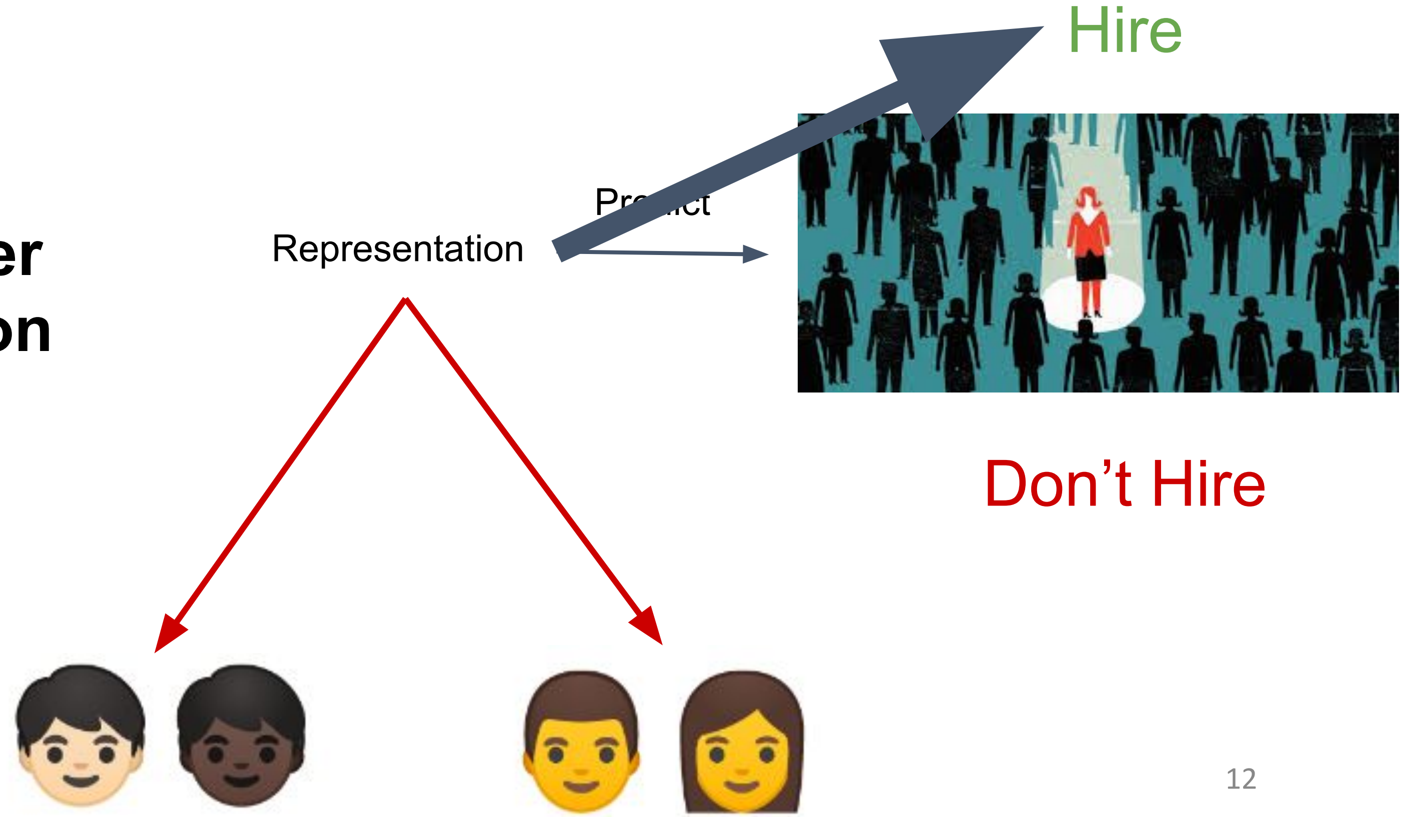

Text classification - Example

In this work:

we do not have access to sensitive tasks like Hiring decisions.

we focus on other tasks, less sensitive

Text classification - Example

Let's predict... EMOJIS

We use DeepMoji.

DeepMoji is a model for predicting Emojis from tweets

Using millions of emoji occurrences to learn any-domain representations for detecting sentiment, emotion and sarcasm

Bjarke Felbo[1], Alan Mislove[2], Anders Søgaard[3], Iyad Rahwan[1], Sune Lehmann[4]

[1]Media Lab, Massachusetts Institute of Technology
[2]College of Computer and Information Science, Northeastern University
[3]Department of Computer Science, University of Copenhagen
[4]DTU Compute, Technical University of Denmark

Text classification - Example

Let's predict... EMOJIS

Deep Moji (Felbo et al., 2017)

I love mom's cooking

I love how you never reply back..

I love cruising with my homies

I love messing with yo mind!!

I love you and now you're just gone..

This is shit

This is the shit

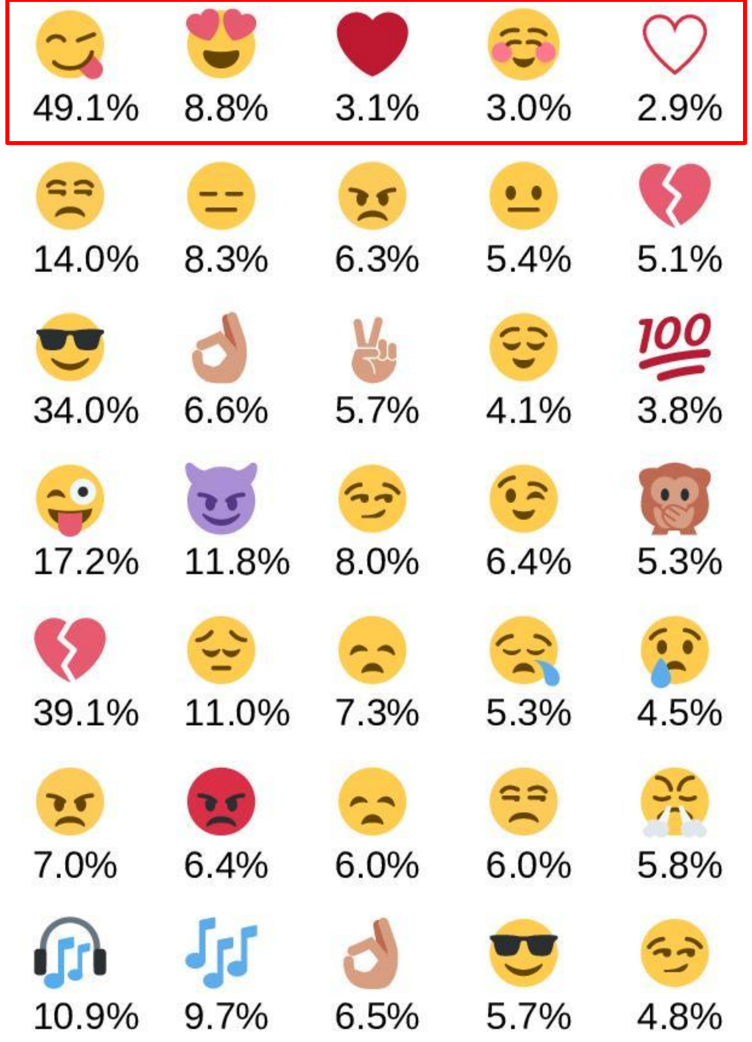

Text classification - Example

Let's predict... EMOJIS

Deep Moji (Felbo et al., 2017)

- DeepMoji is a strong and expressive model
- It also create powerful **representations**

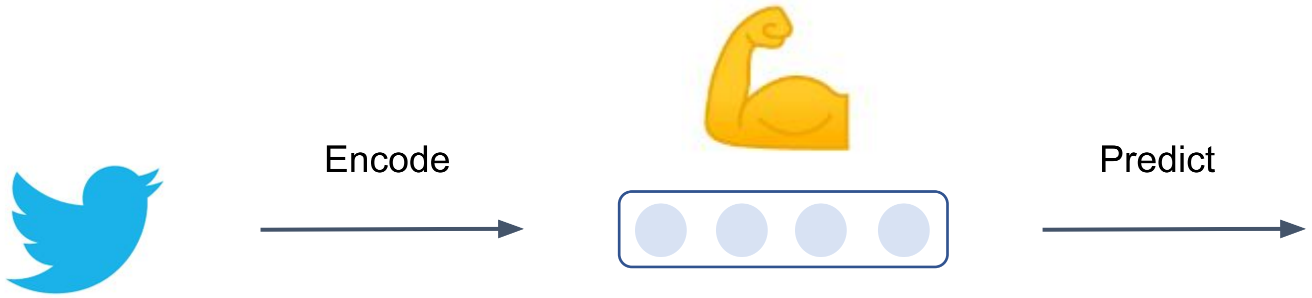

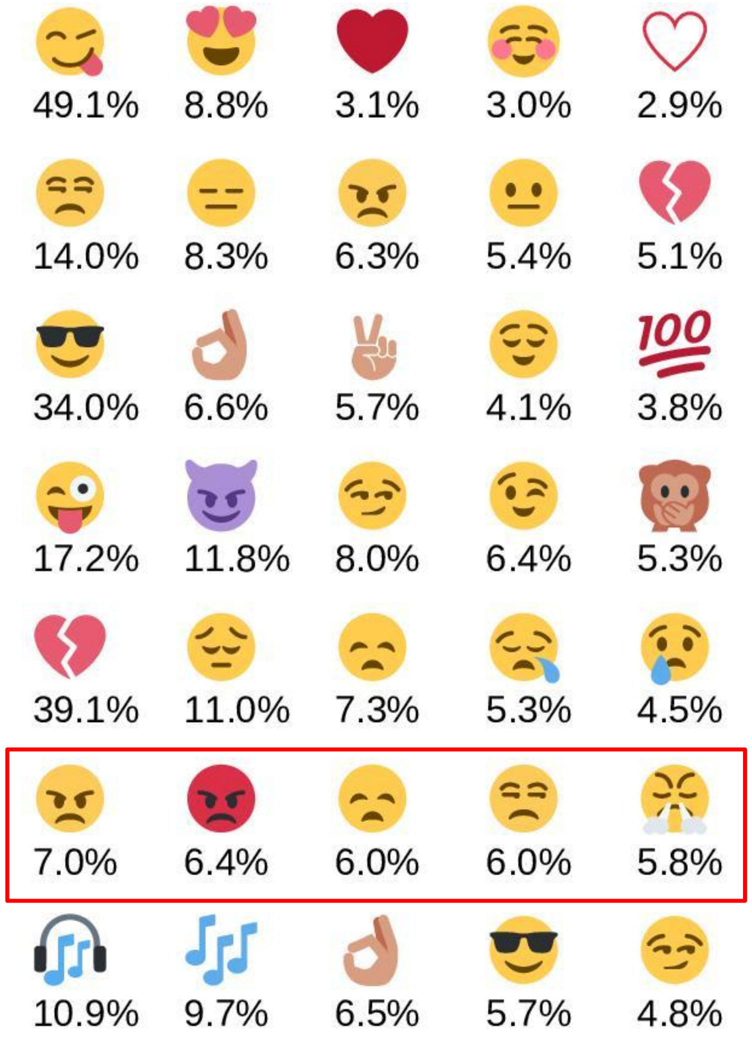

Text classification - Example

Let's predict... EMOJIS

Deep Moji (Felbo et al., 2017)

- DeepMoji is a strong and expressive model
- It also create powerful **representations**

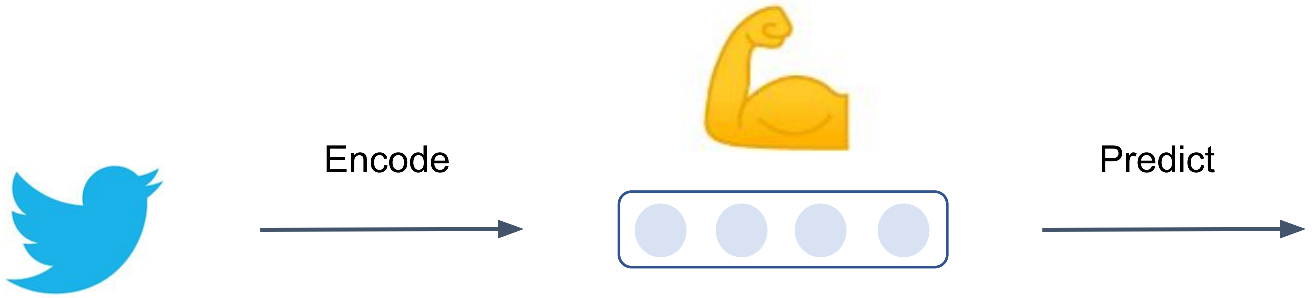

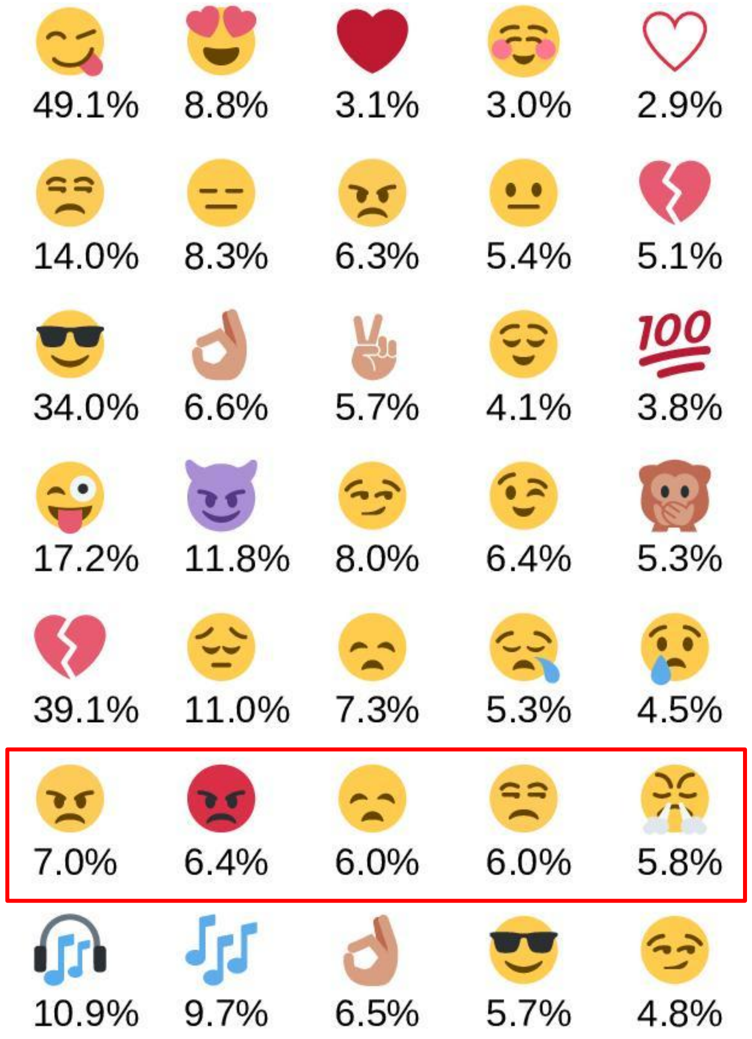

- Achieved several SOTA results on text classification

Text classification - Example

Let's predict... EMOJIS

Does this representation also contain information on sensitive attributes?

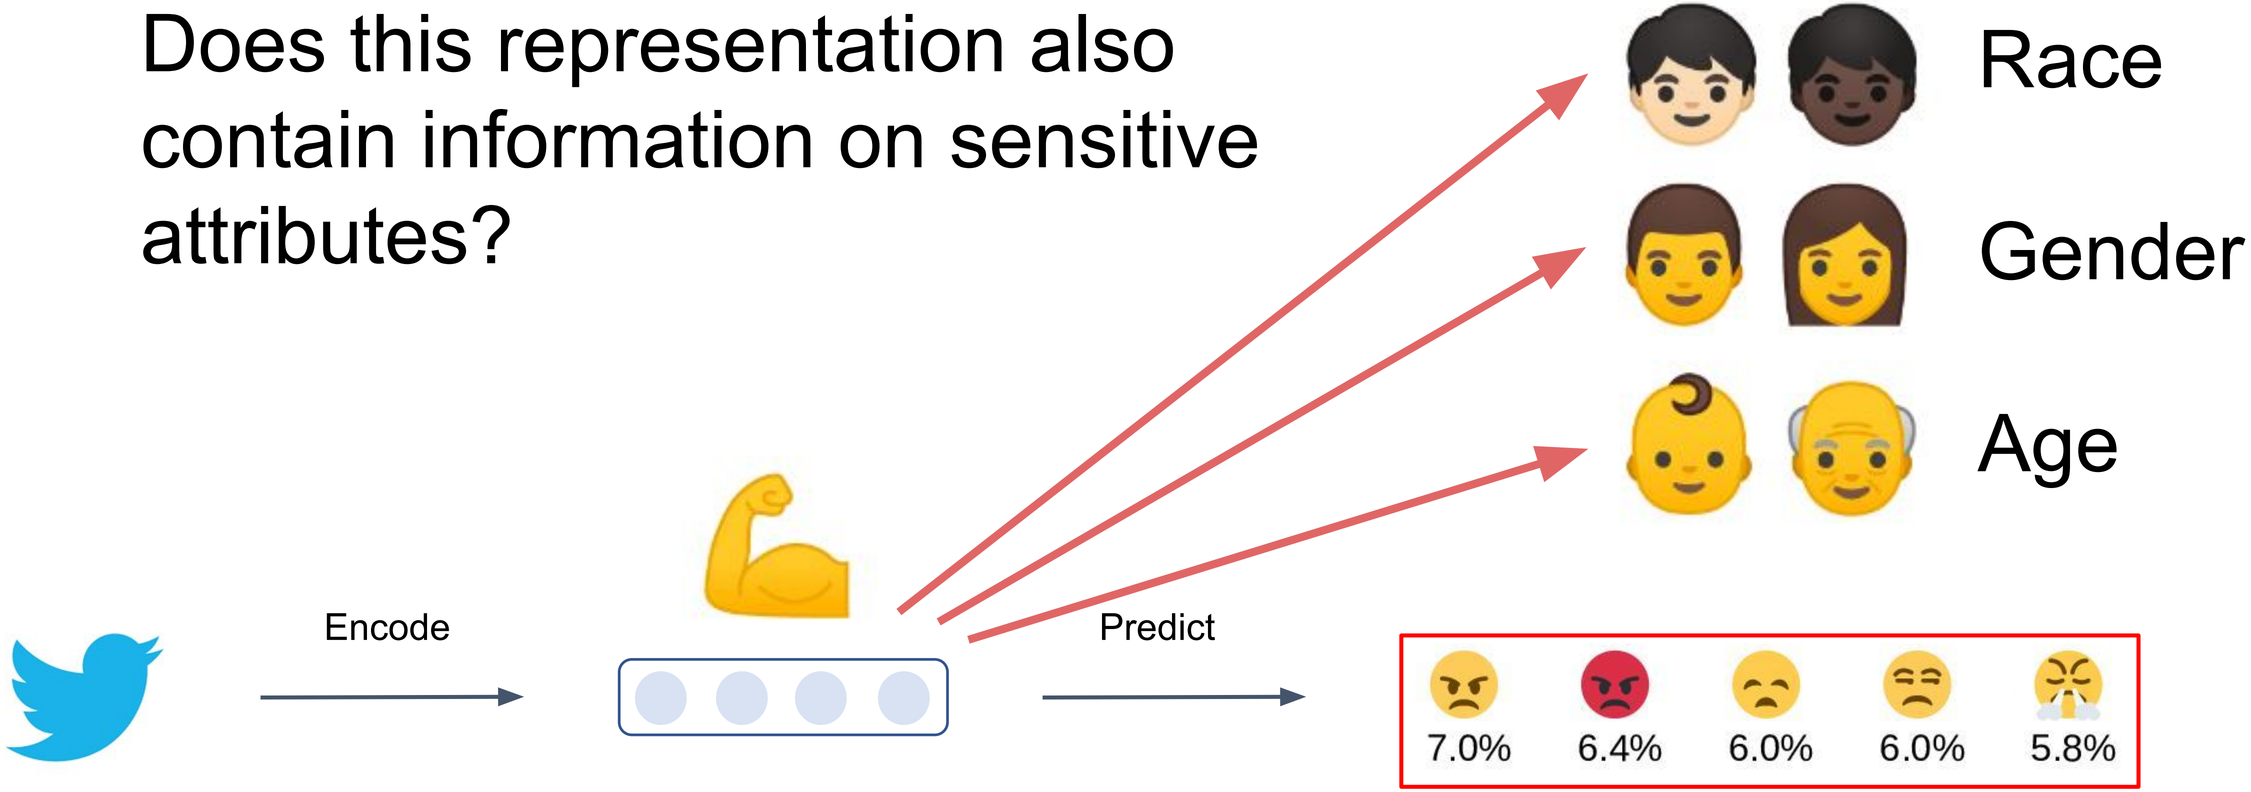

Setup

Task
(Emojis)

We take the representation that predict Emojis

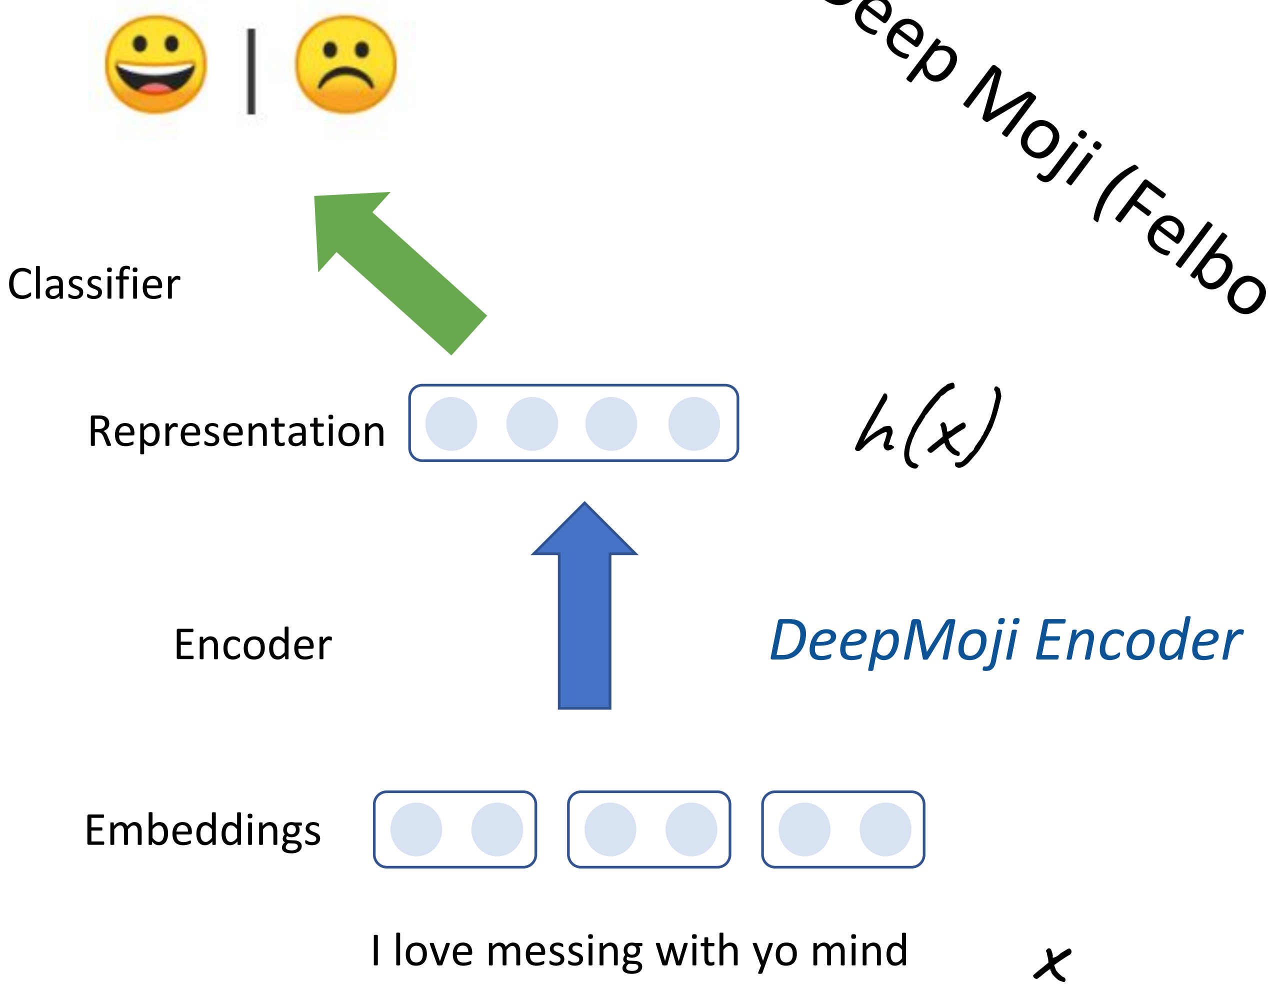

Deep Moji (Felbo et al., 2017)

Setup

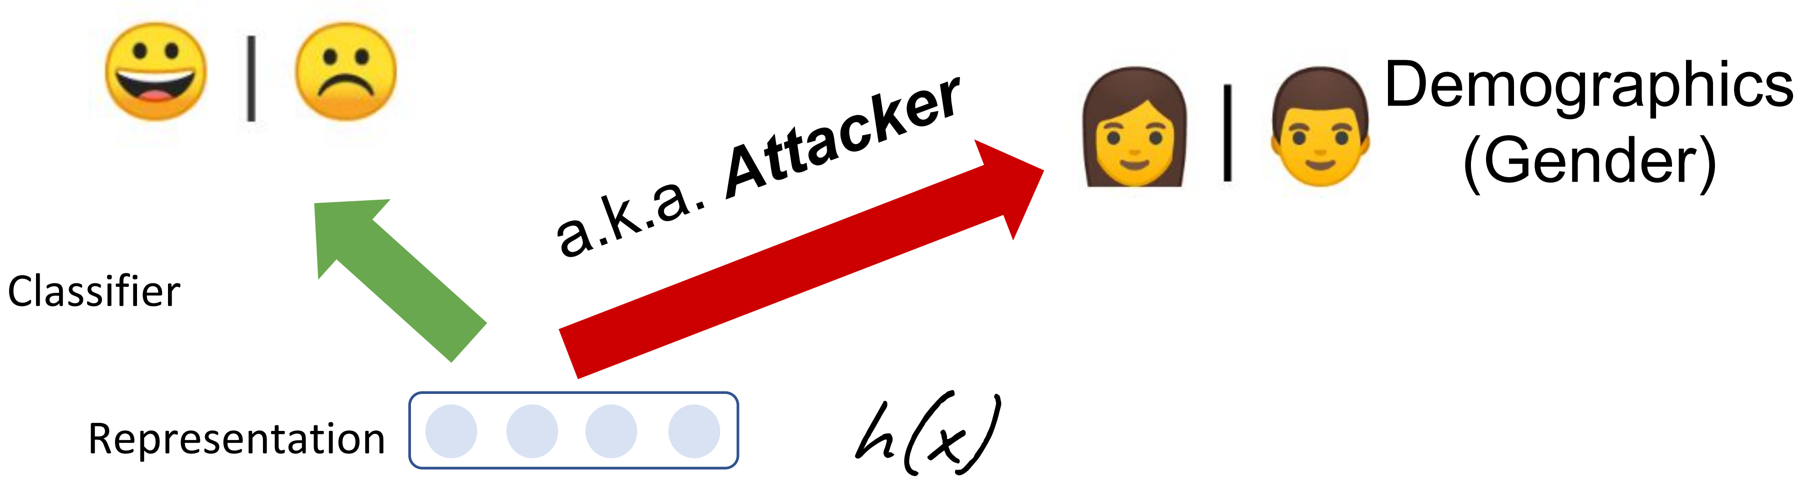

We take the representation that predict Emojis

And use them to predict demographics.

We define:
<u>leakage</u> = score above a random guess an “**Attacker**” achieves

Text Leakage – Case Study

- We use DeepMoji encoder, to encode tweets, from 3 datasets, all binary and balanced

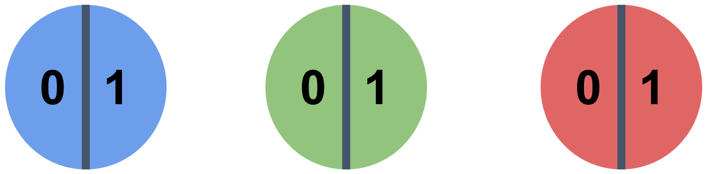

- Each dataset is tied to a different demographic label

- We then train Attackers to predict these attributes

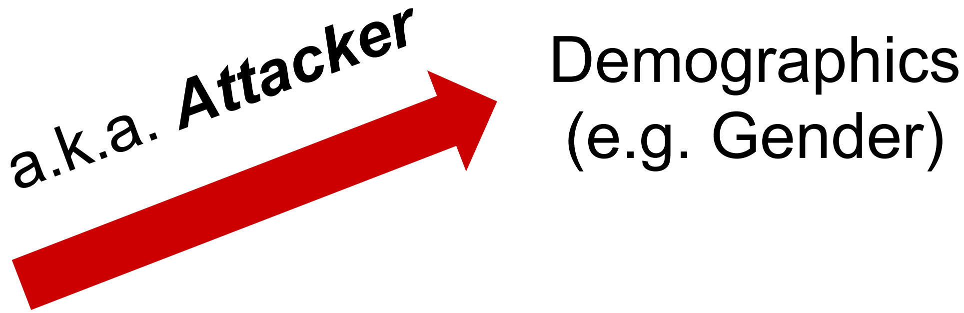

Text Leakage – Case Study

The dev-set scores above chance level are quite high

Big Surprise?

Not really.
This is the core idea in **Transfer-Learning**.
We've seen its benefits in pretrained embeddings, language models etc.

Random Guess

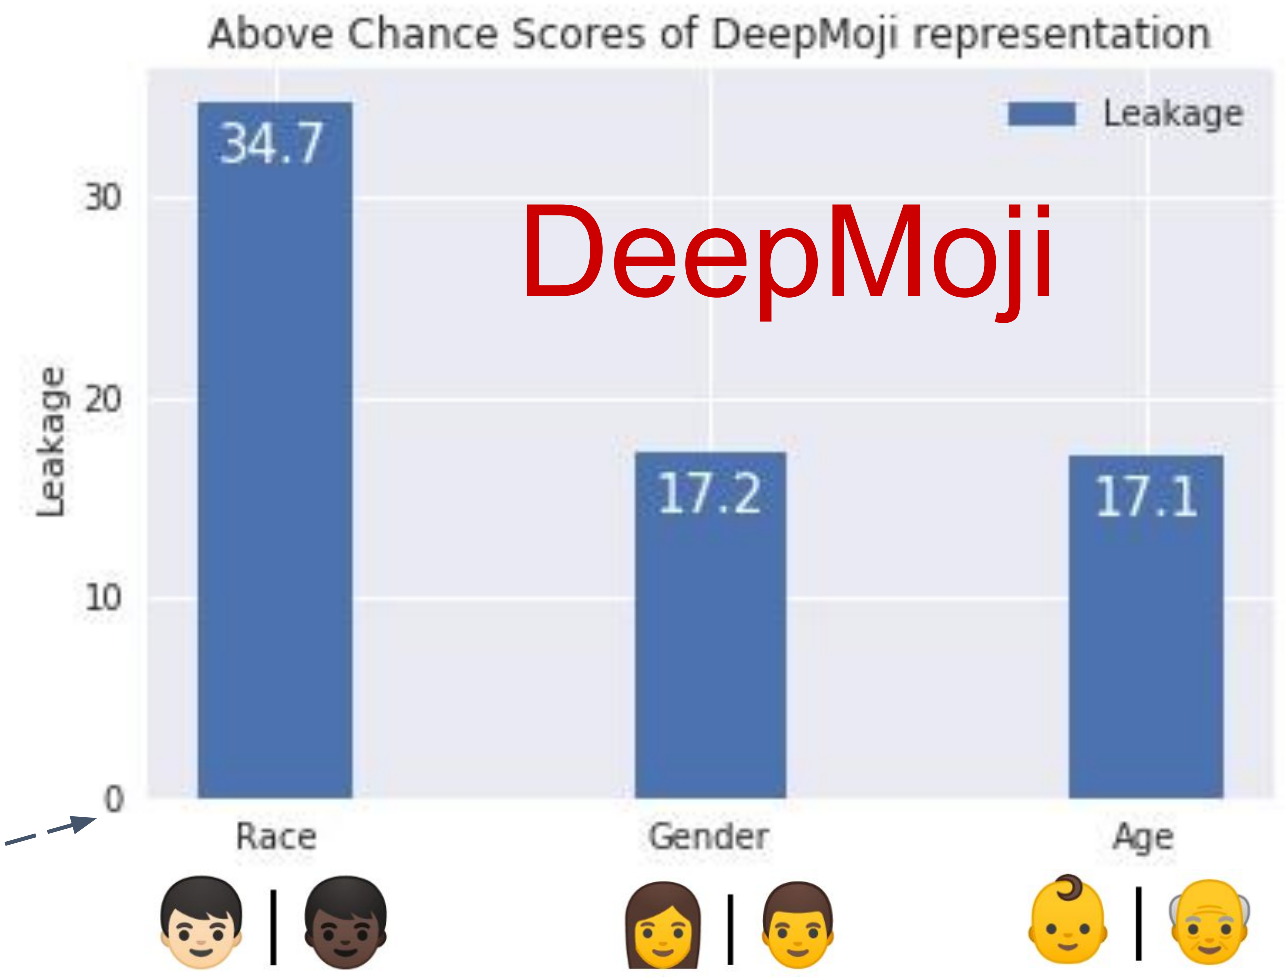

Text Leakage – Case Study

- Why do we get this major “help” in predicting other attributes than those we trained for?
- One option is the correlation between attributes in the data.

Fair enough. Let’s control for it.

Controlled Setup

New setup

- We use Twitter data
- We focus on sentiment prediction, emoji based 😀 | ☹️
- With *Race*, *Gender* and *Age* as protected attributes

👦🏻 | 👦🏿 👩 | 👨 👶 | 👴

Blodgett et al., 2016 *Rangel et al., 2016* *Rangel et al., 2016*

New setup

Balanced Dataset

Demographics

50% Male *50% Female*

Task (Sentiment)

50% Positive

50% Negative

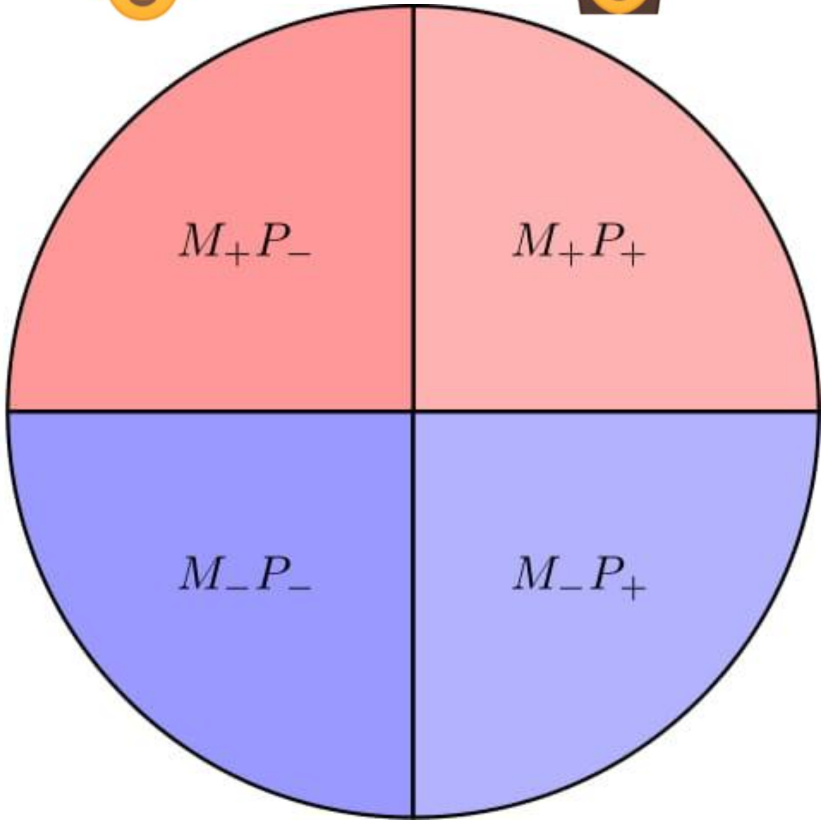

Balanced Training

Training our own encoder on the balanced datasets

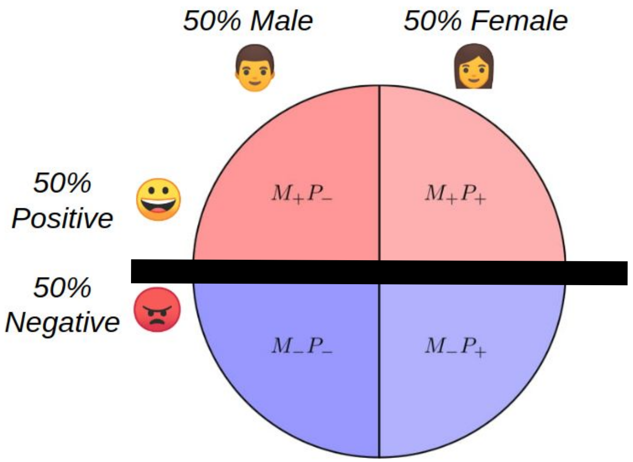

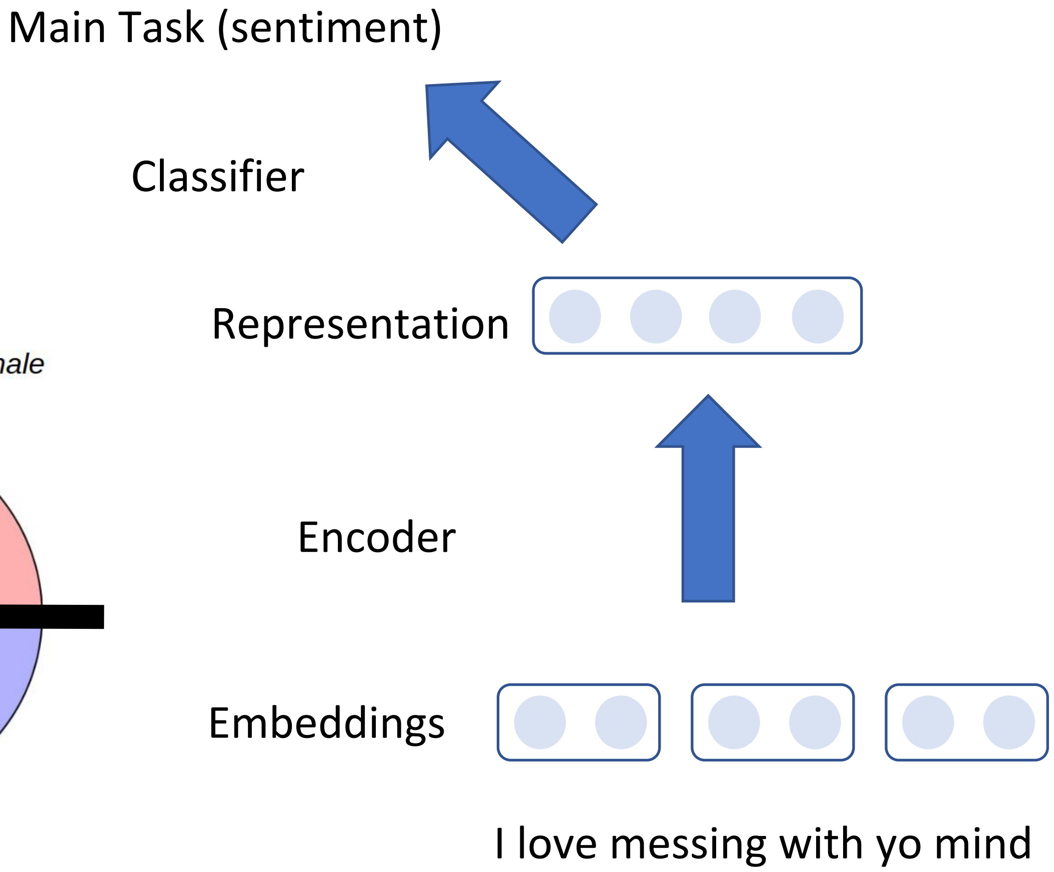

Balanced Training

And using the Attacker to check for leakage

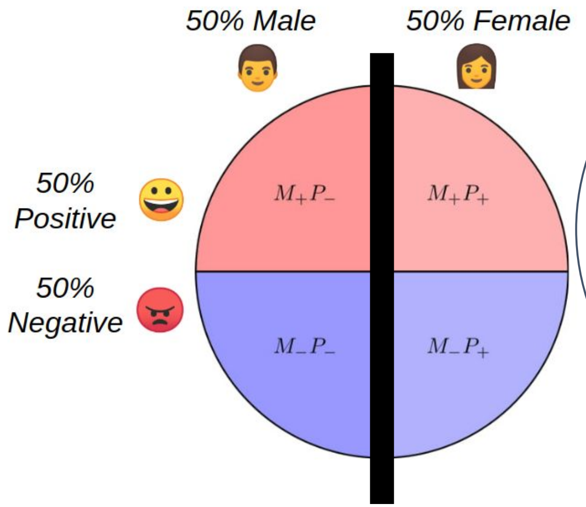

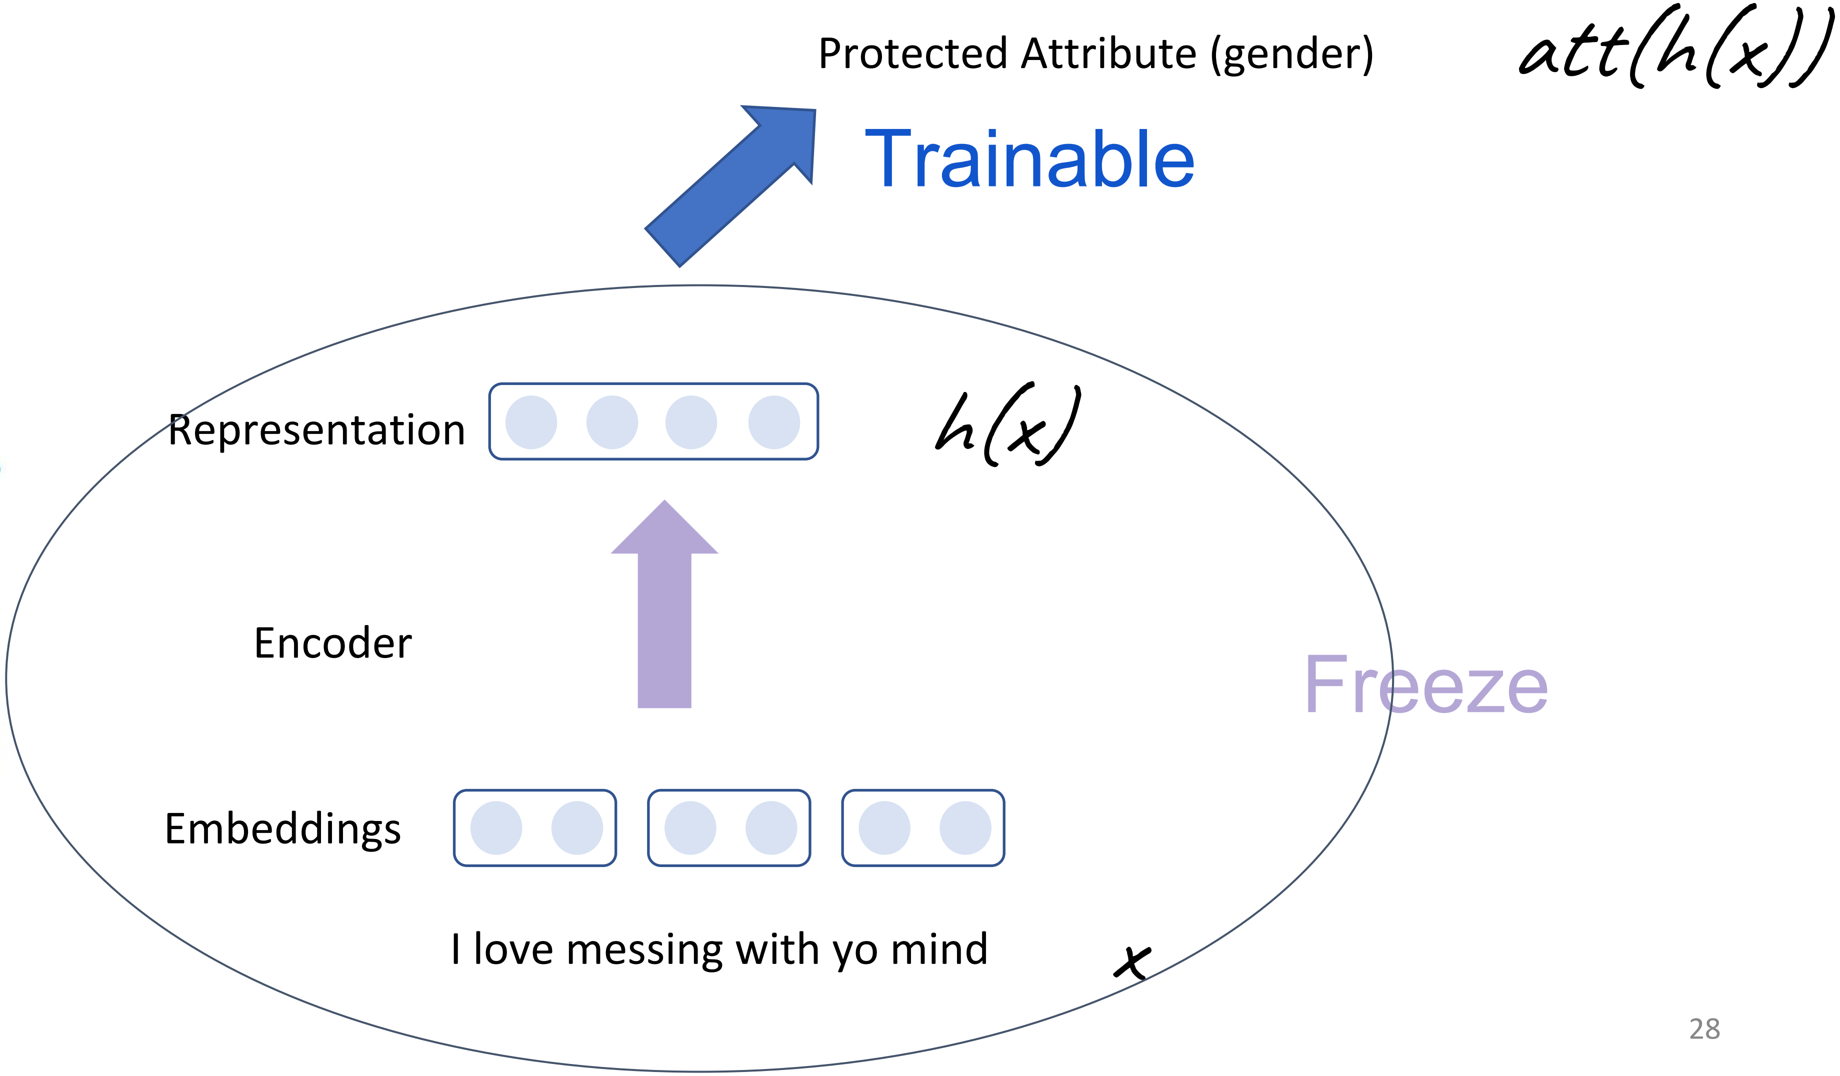

Balanced Training - Leakage

We wanted to see something like this:

But instead...

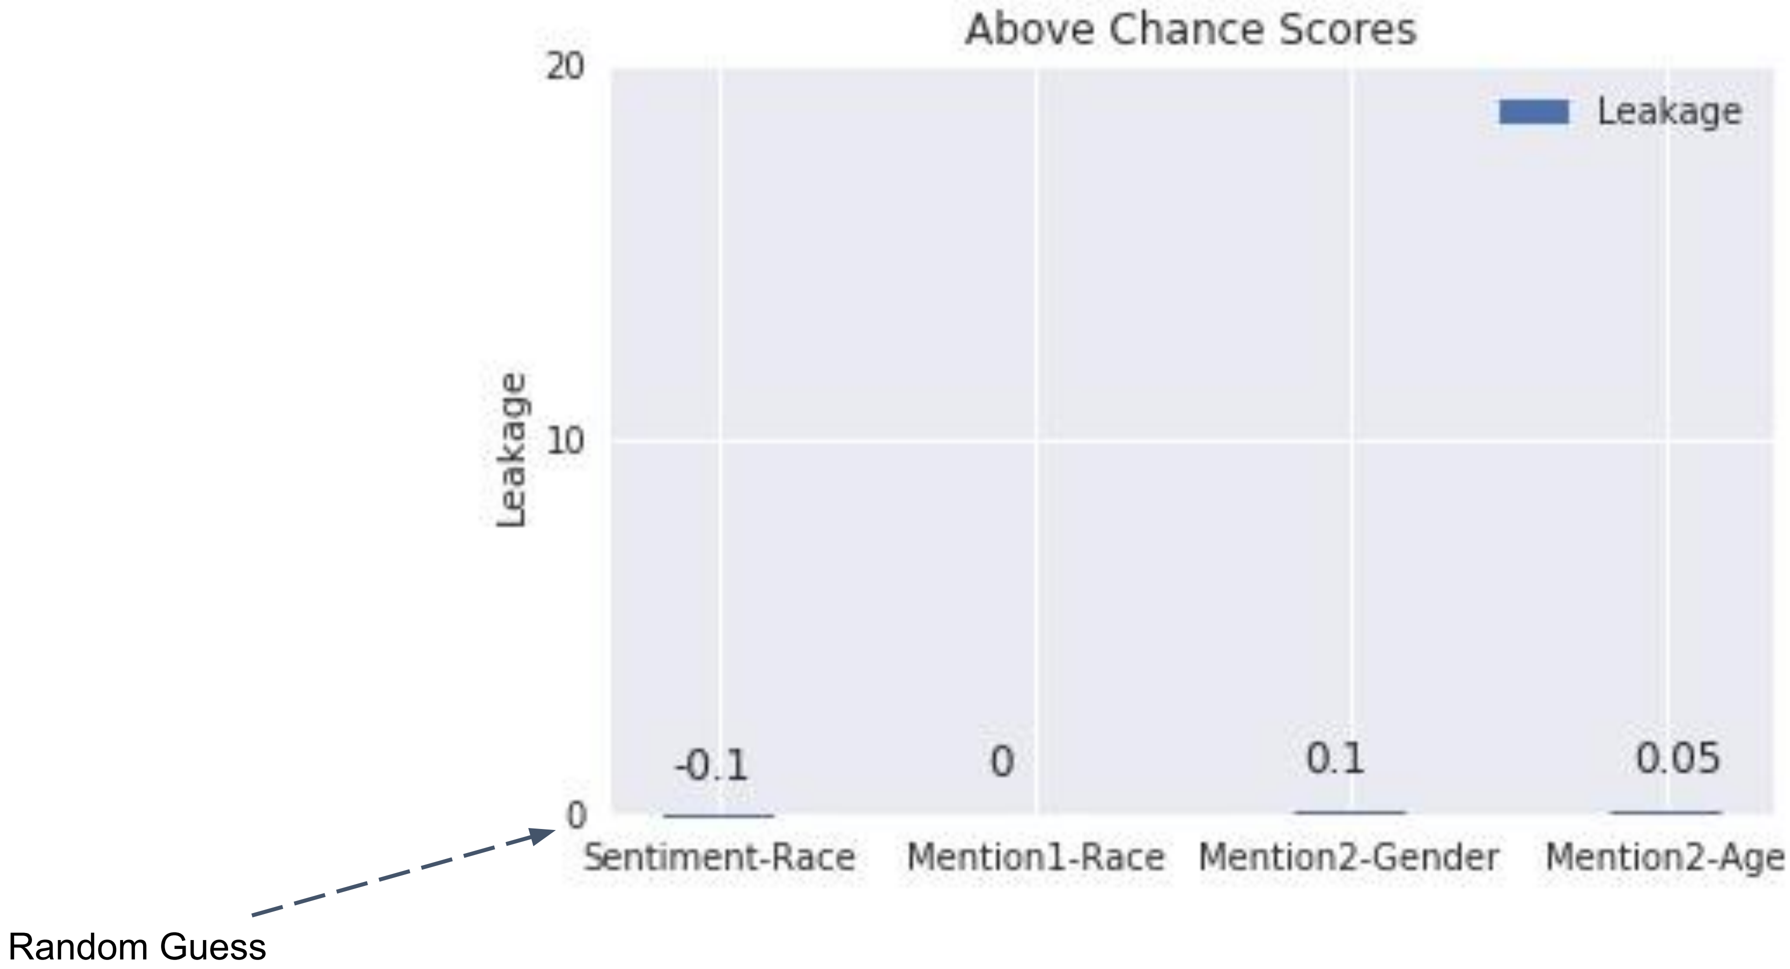

Balanced Training - Leakage

The Attacker manages to extract a substantial amount of sensitive information

Even in a balanced setup, leakage exists

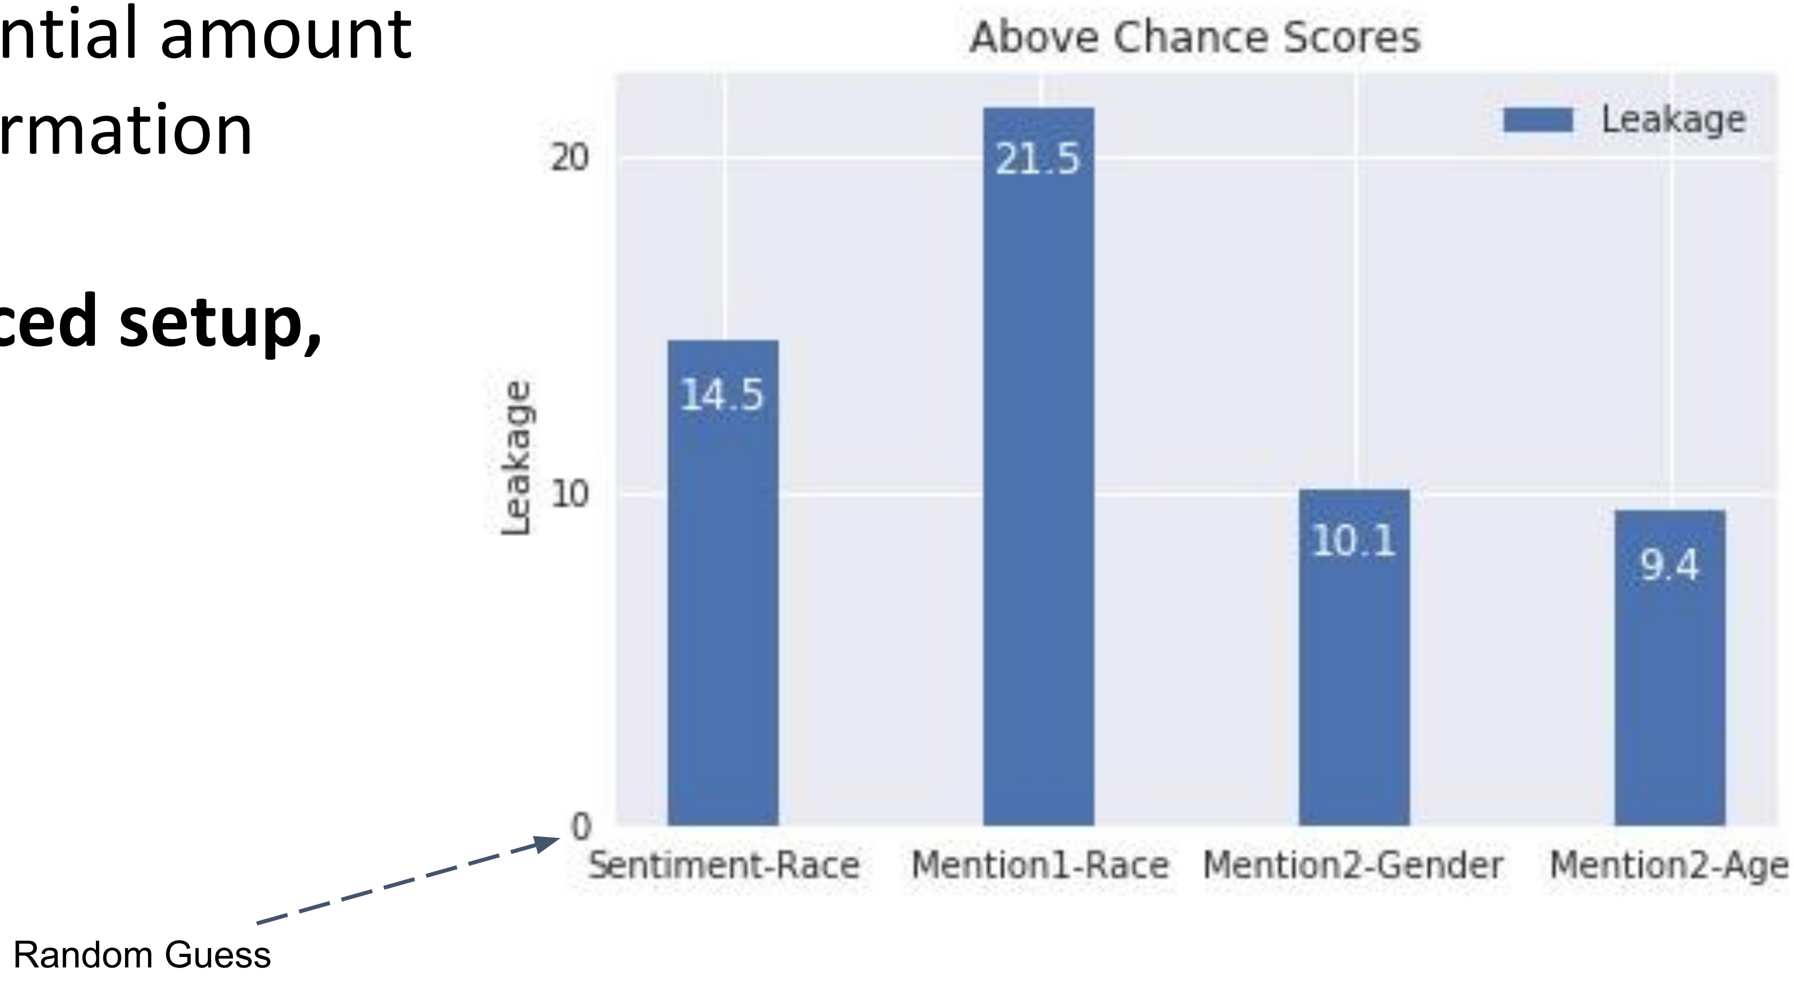

Our objective

- Create a representation which:
  - Is predictive of the main task (e.g. sentiment)

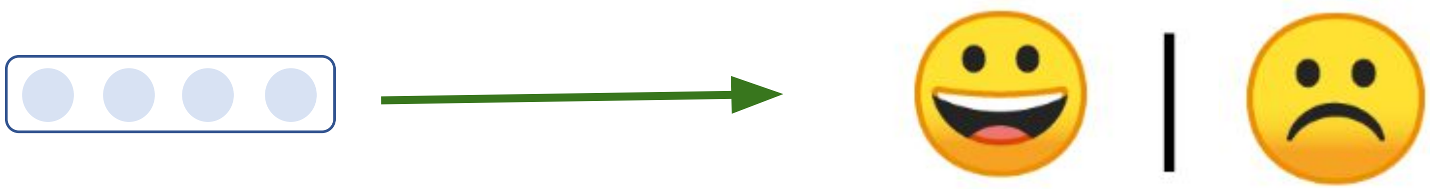

Our objective

- Create a representation which:
  - Is predictive of the main task (e.g. sentiment)

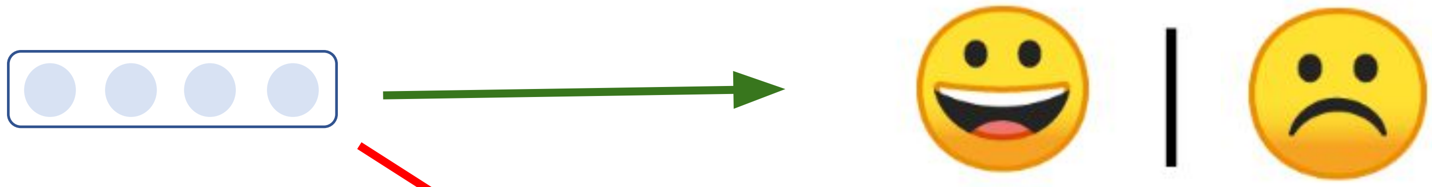

and

  - Is **not** predictive of protected attribute (e.g. gender, race)

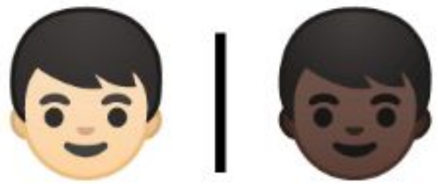

Our objective

- Interesting technical problem – How to **unlearn** something?
- Interesting technical problem – Can we **unlearn** something?

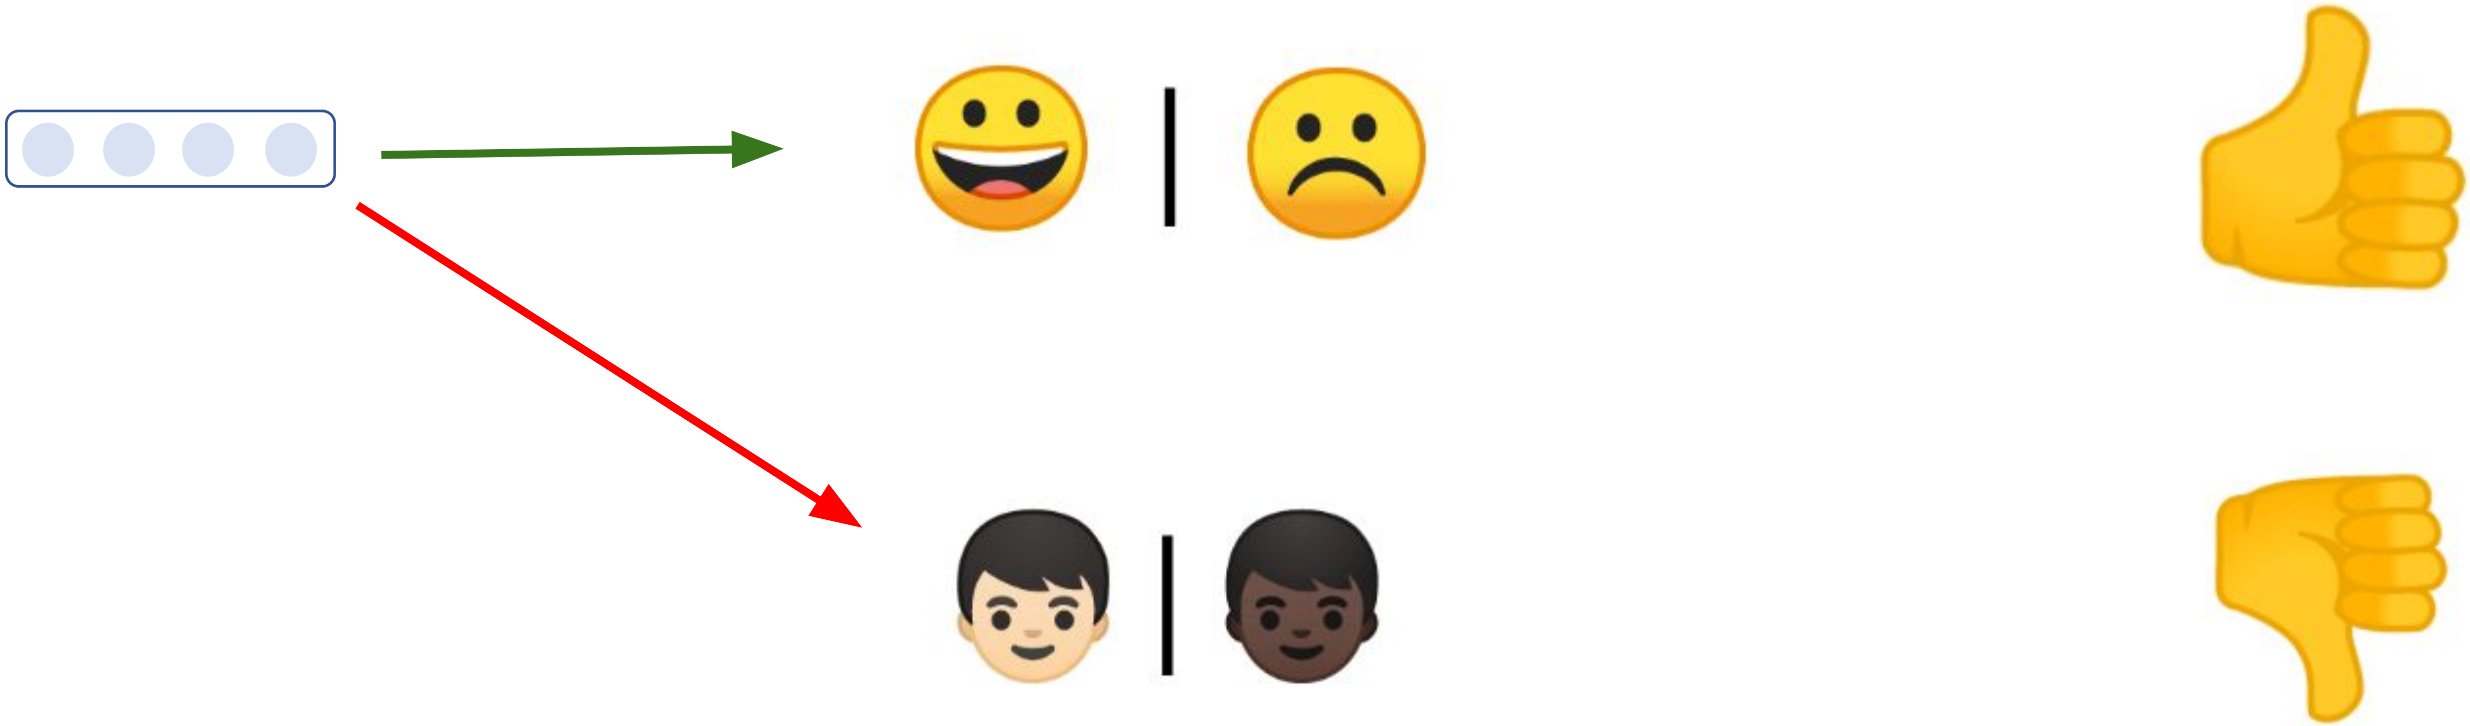

Actively Reducing Leakage

Adversarial Setup (Ganin and Lempitsky, 2015)

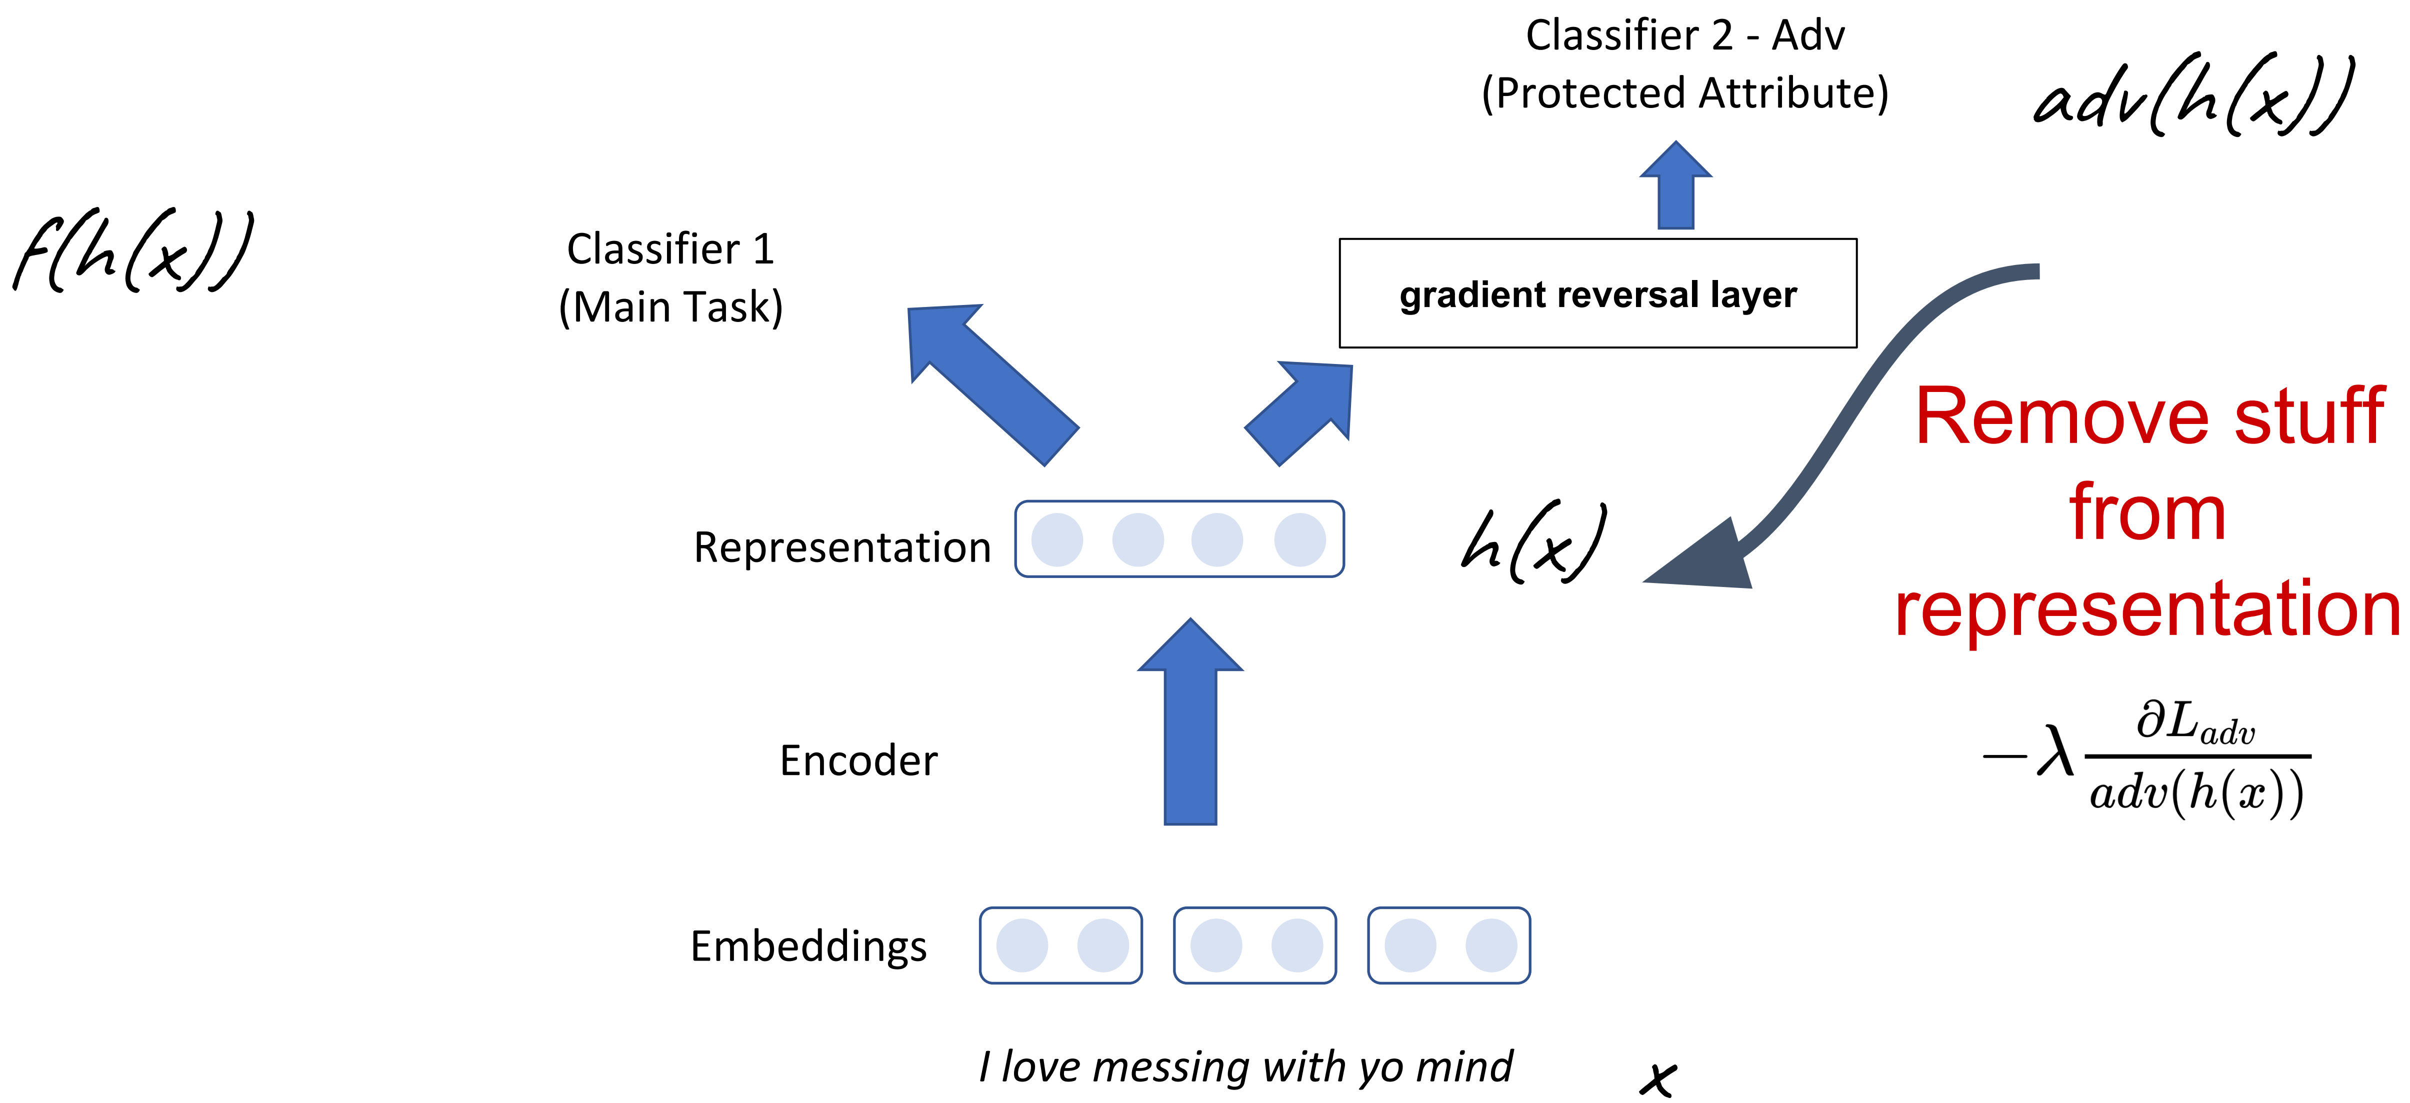

$$-\lambda \frac{\partial L_{adv}}{adv(h(x))}$$

Does it work?

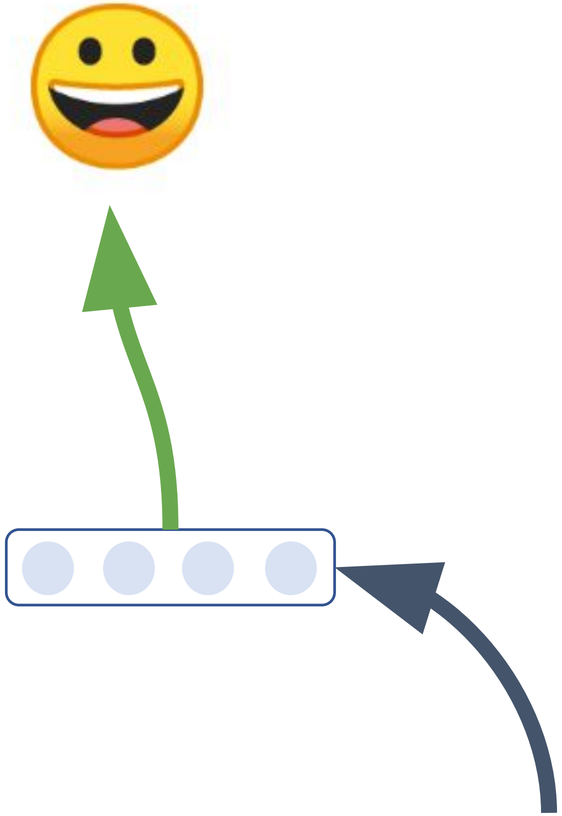

Successfully predicting sentiment

"I love mom's cooking"

Does it work?

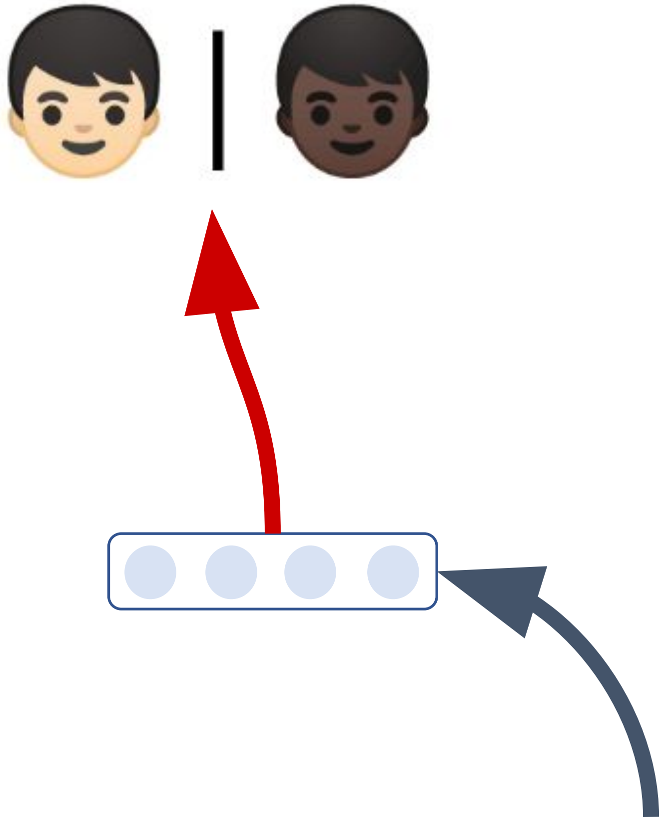

Successfully removed demographics?

"I love mom's cooking"

Does it work?

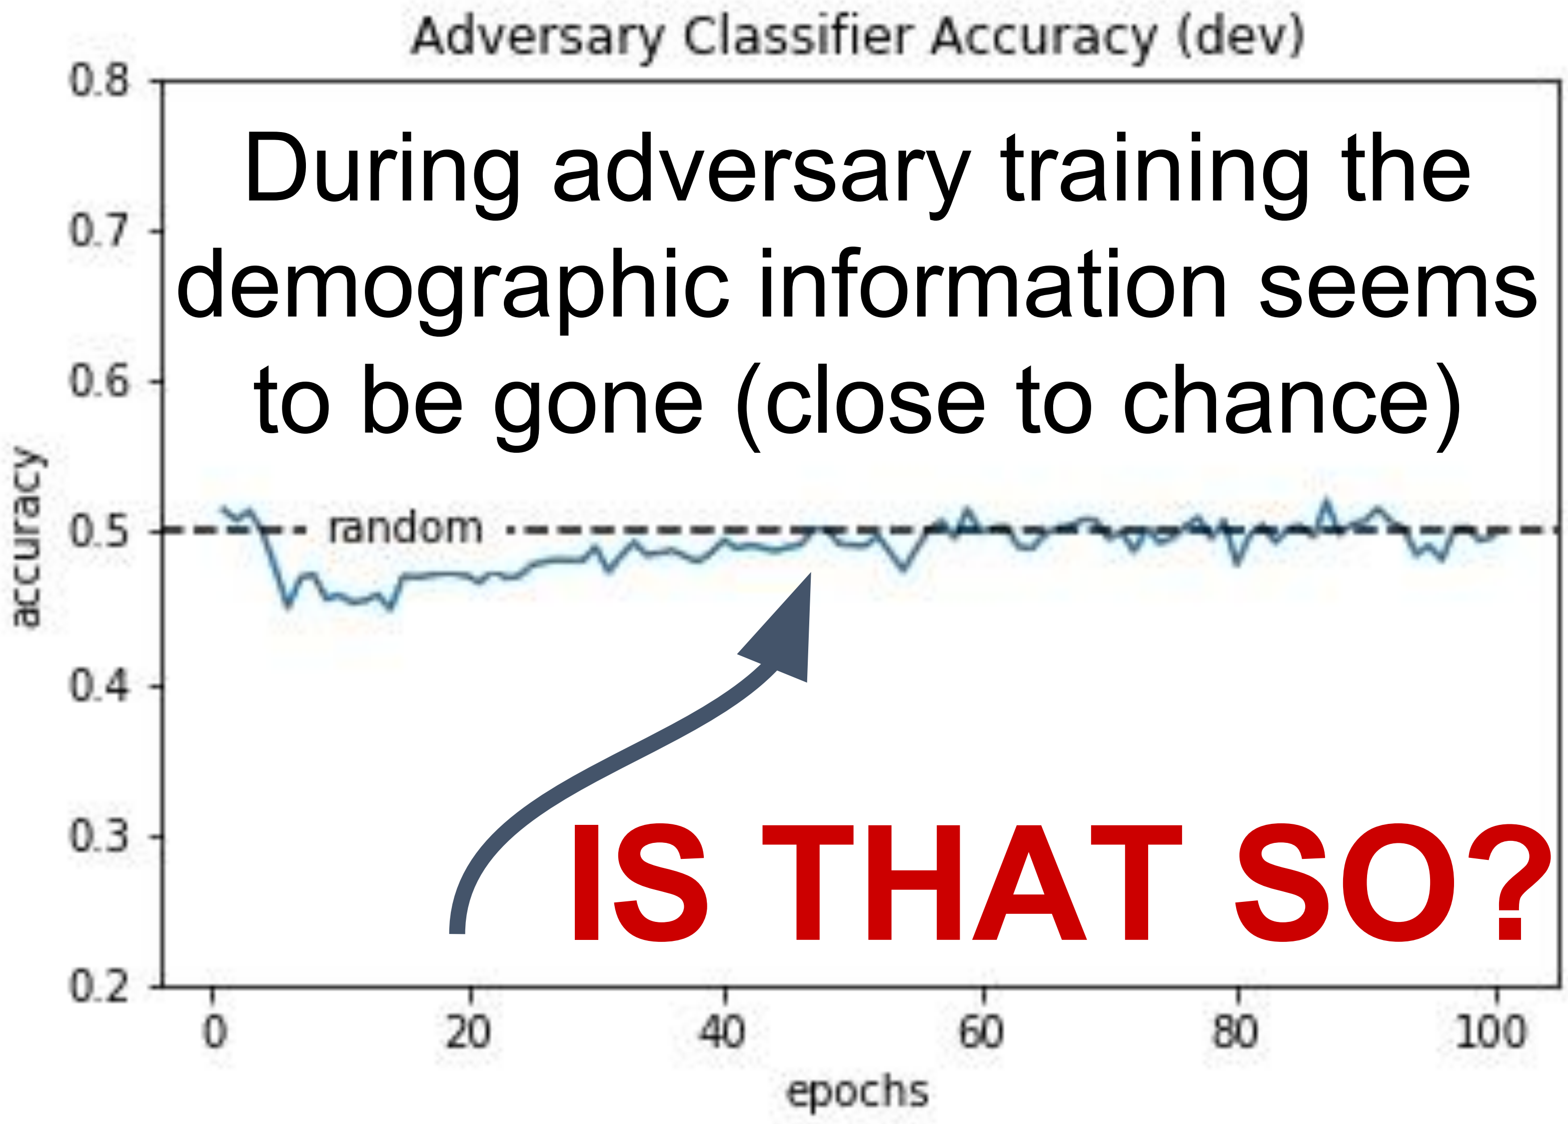

During adversary training the demographic information seems to be gone (close to chance)

IS THAT SO?

Does it work? Not so quickly...

When training the Attacker

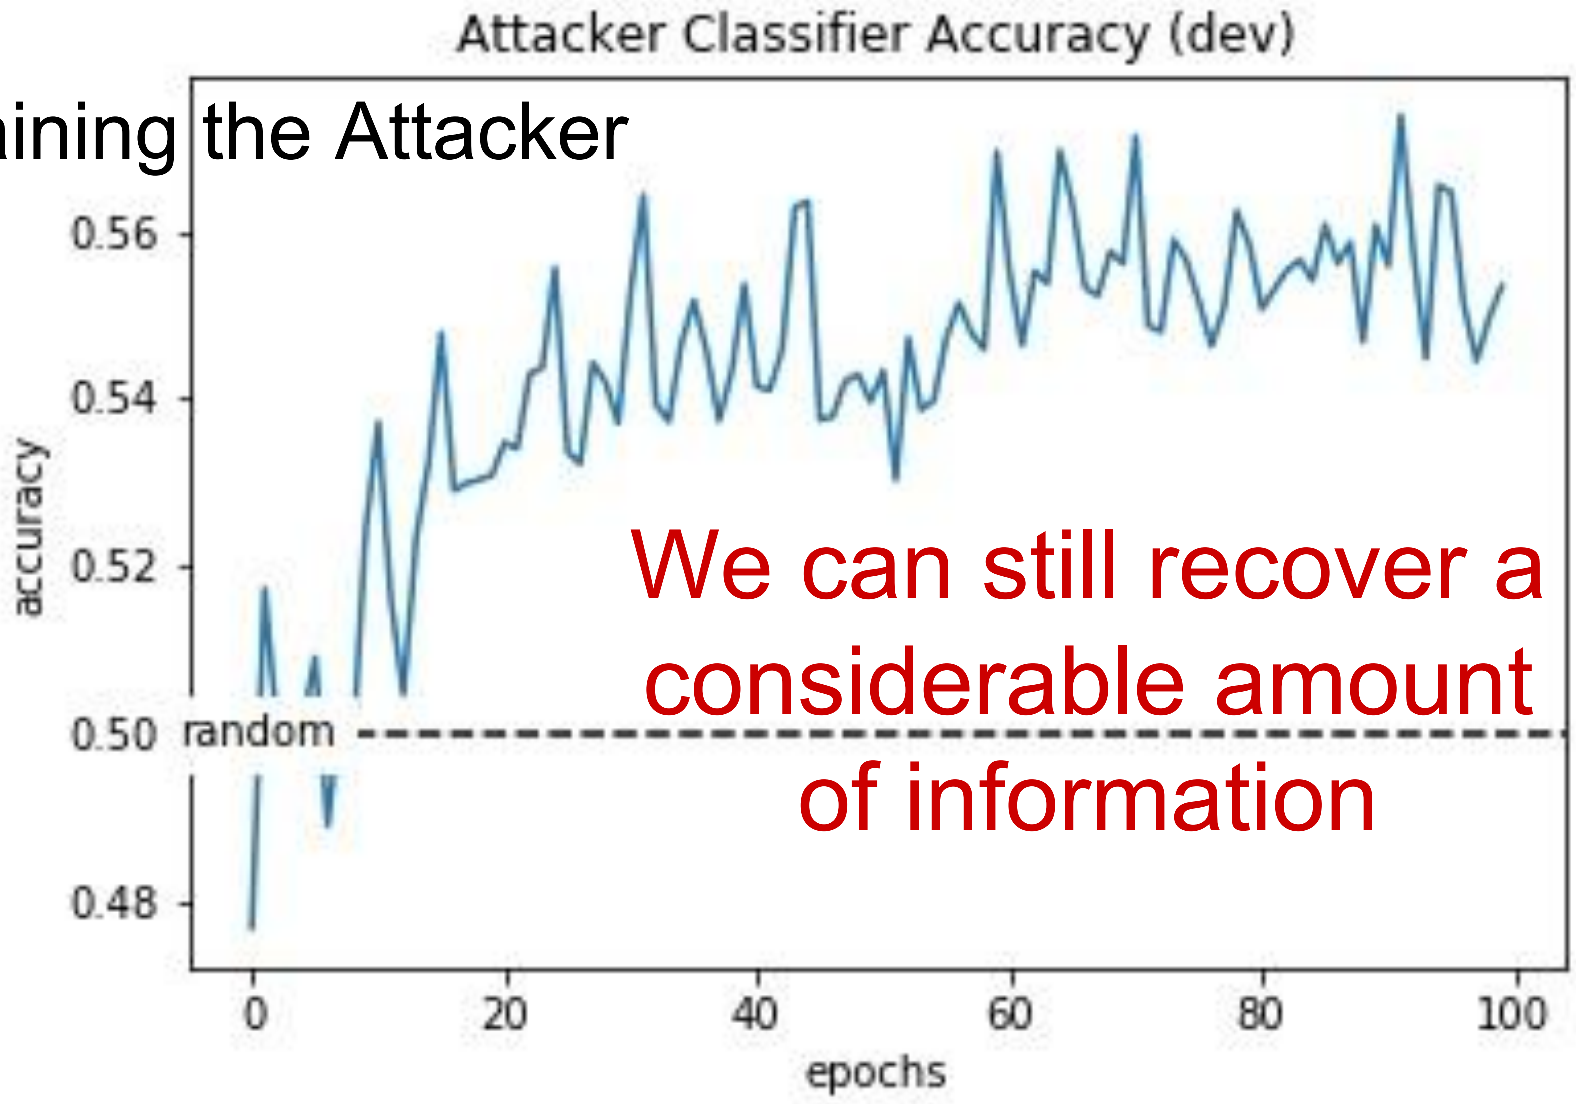

We can still recover a considerable amount of information

Does it work? Not so quickly...

Consistent across tasks and protected attributes

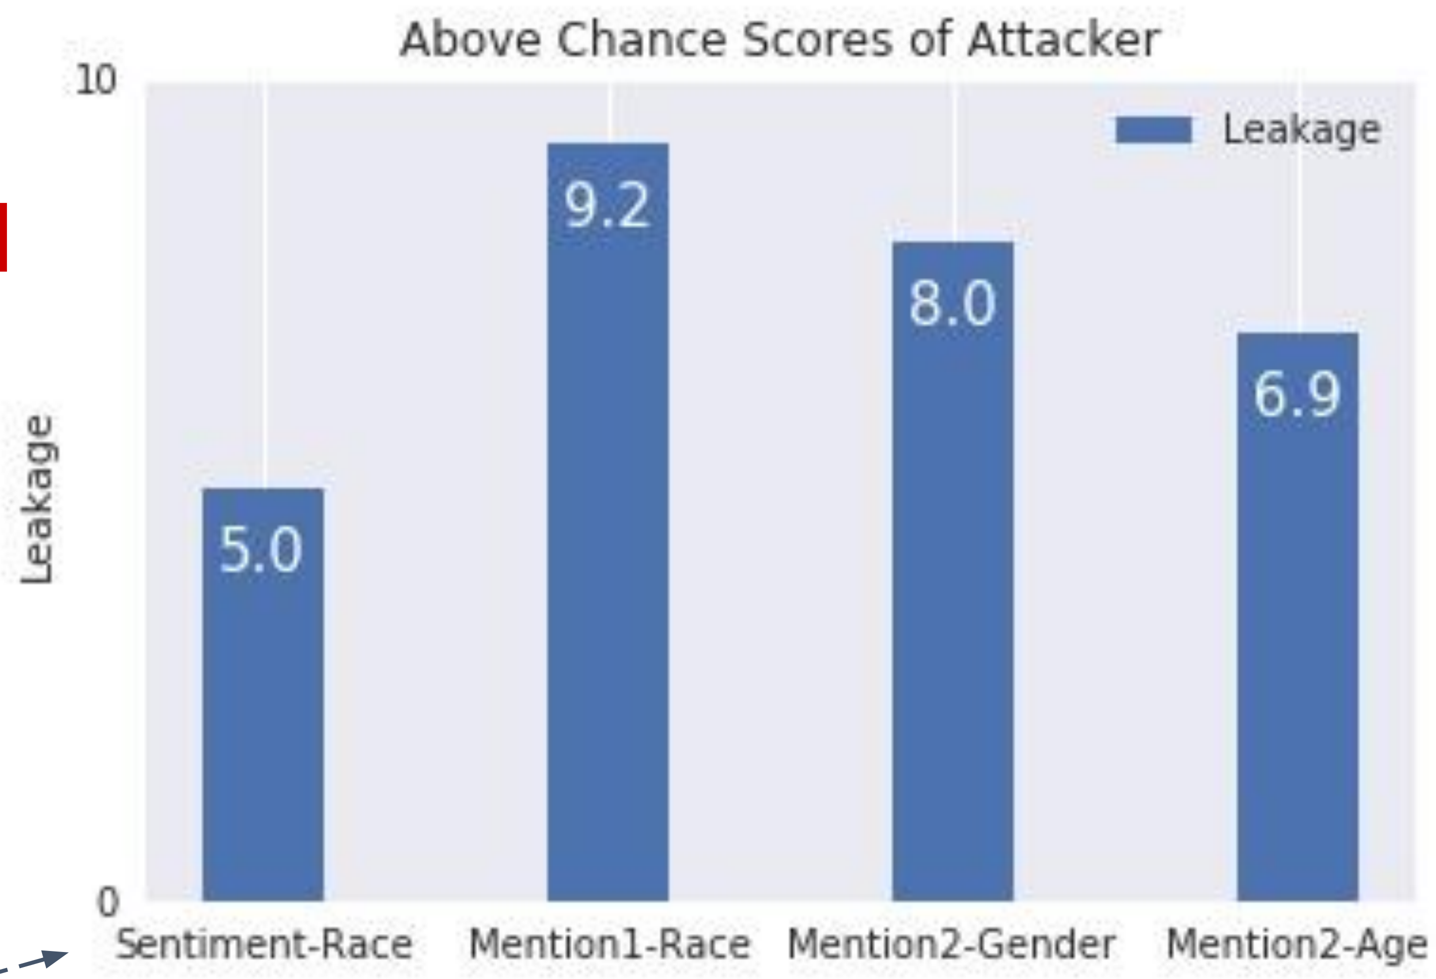

Random Guess

Does it work? more or less

Well, the adversarial method does help.
But not enough

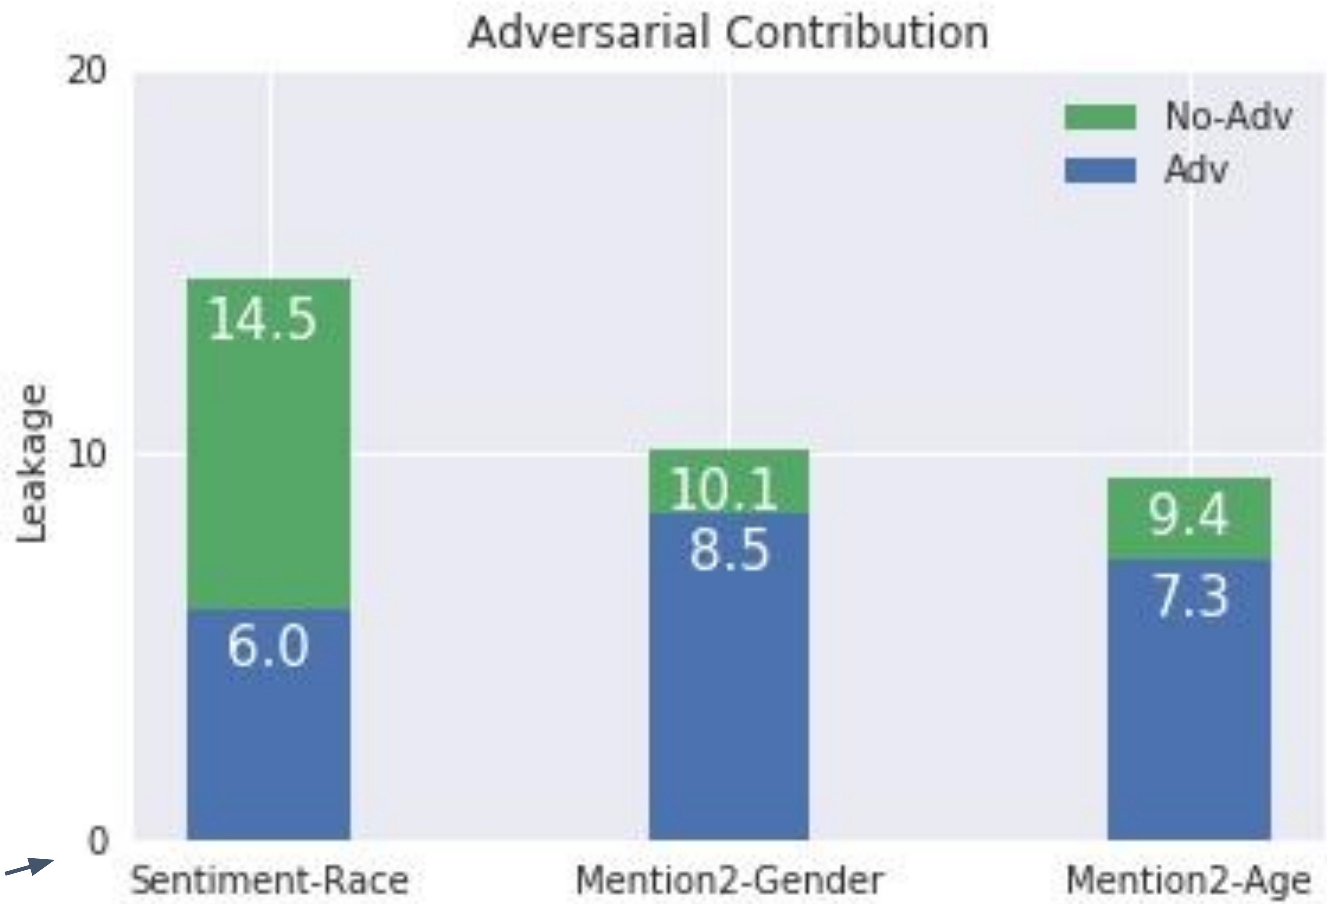

Random Guess

While effective during training, in test time, the adversarial do not remove all the protected information

Stronger, Better, Bigger???

Can we make stronger adversaries?

Stronger, Better, Bigger???

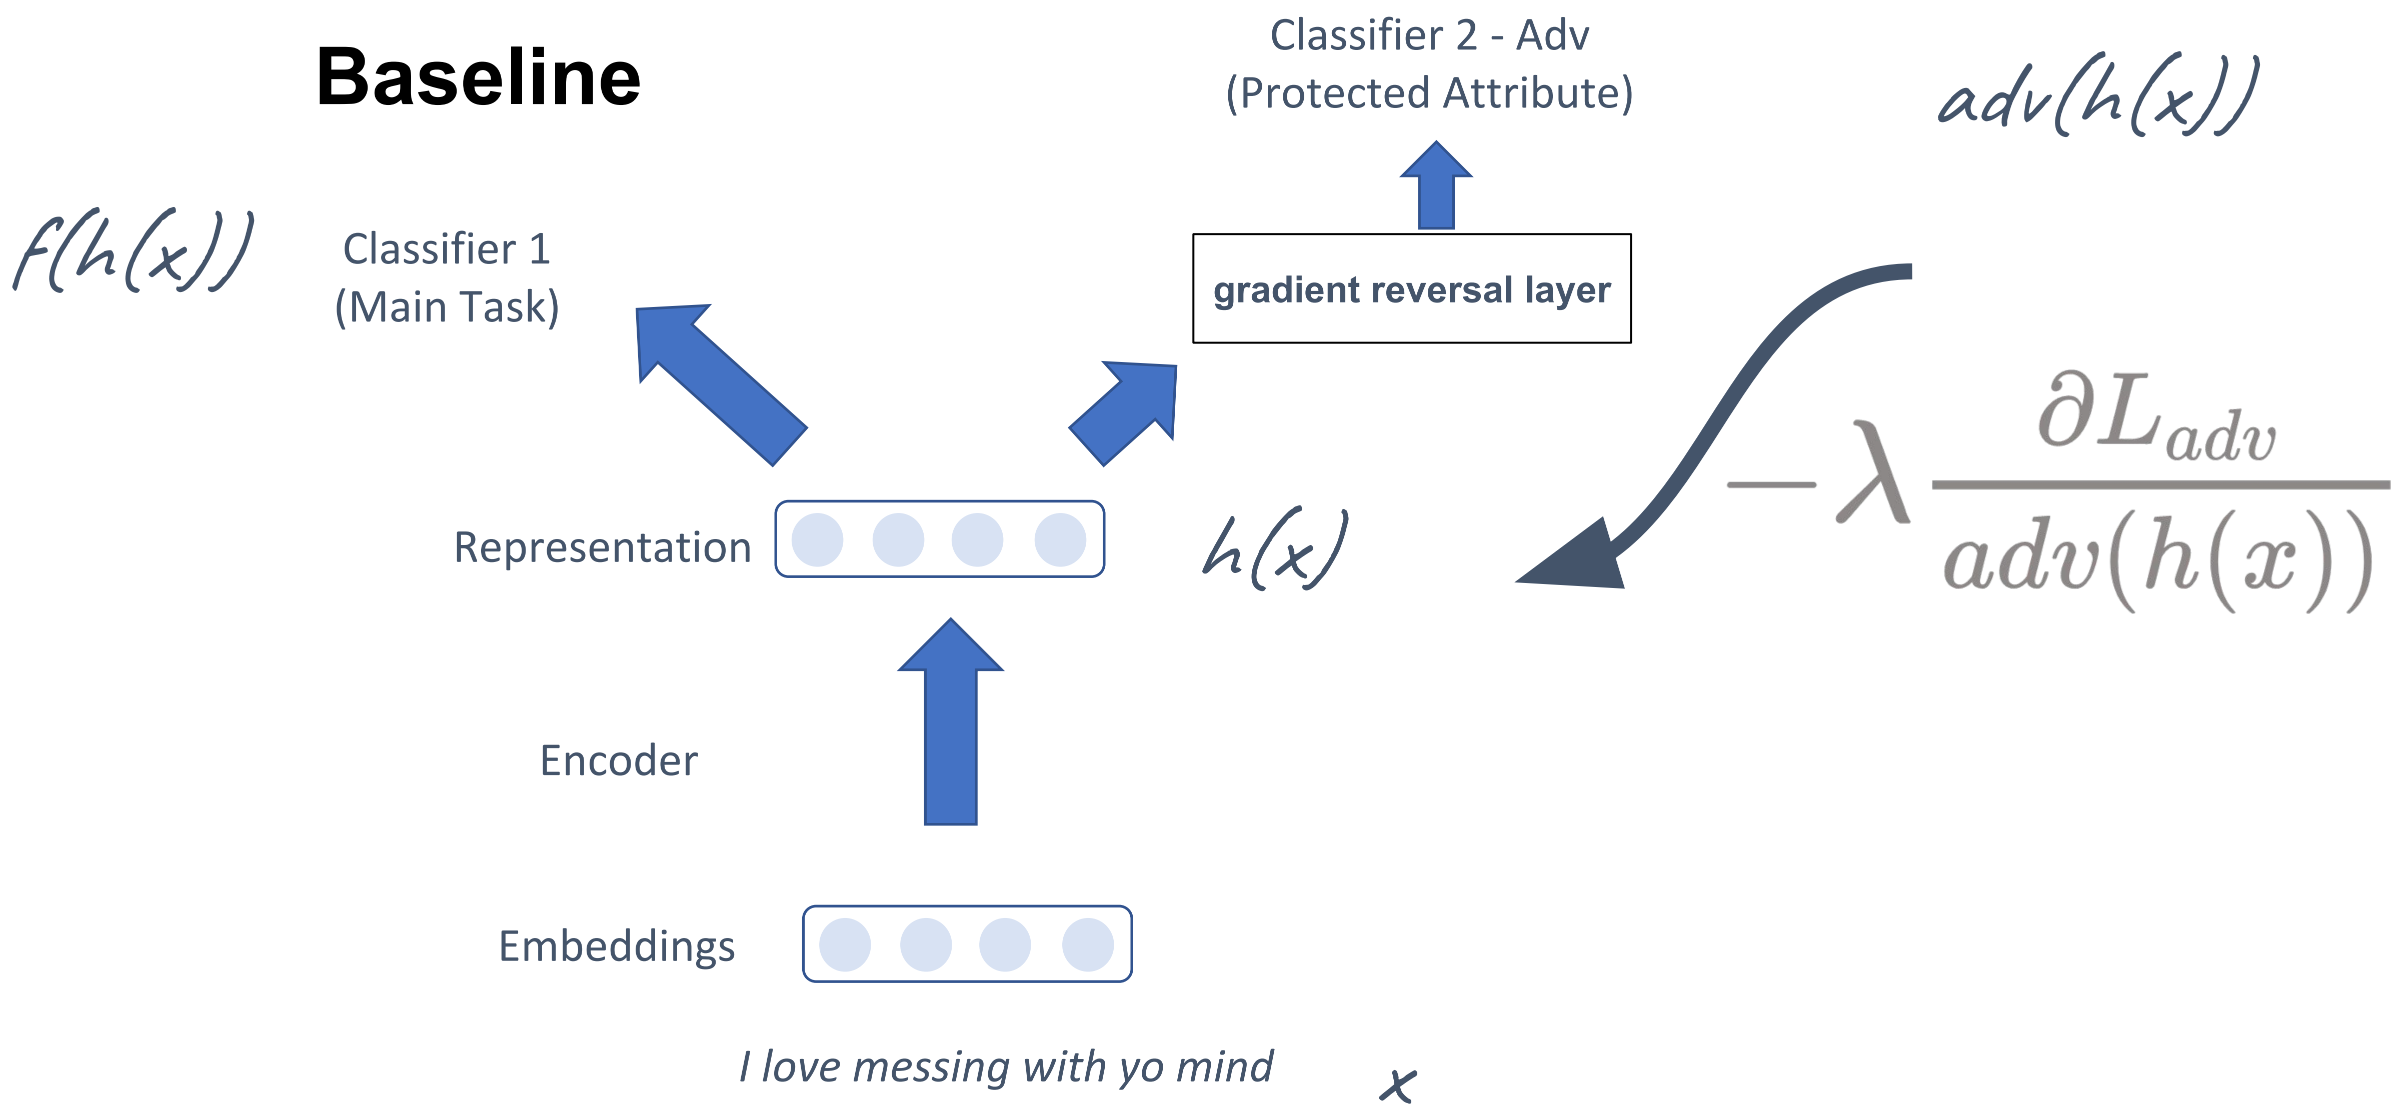

Stronger, Better, Bigger???

More Parameters!

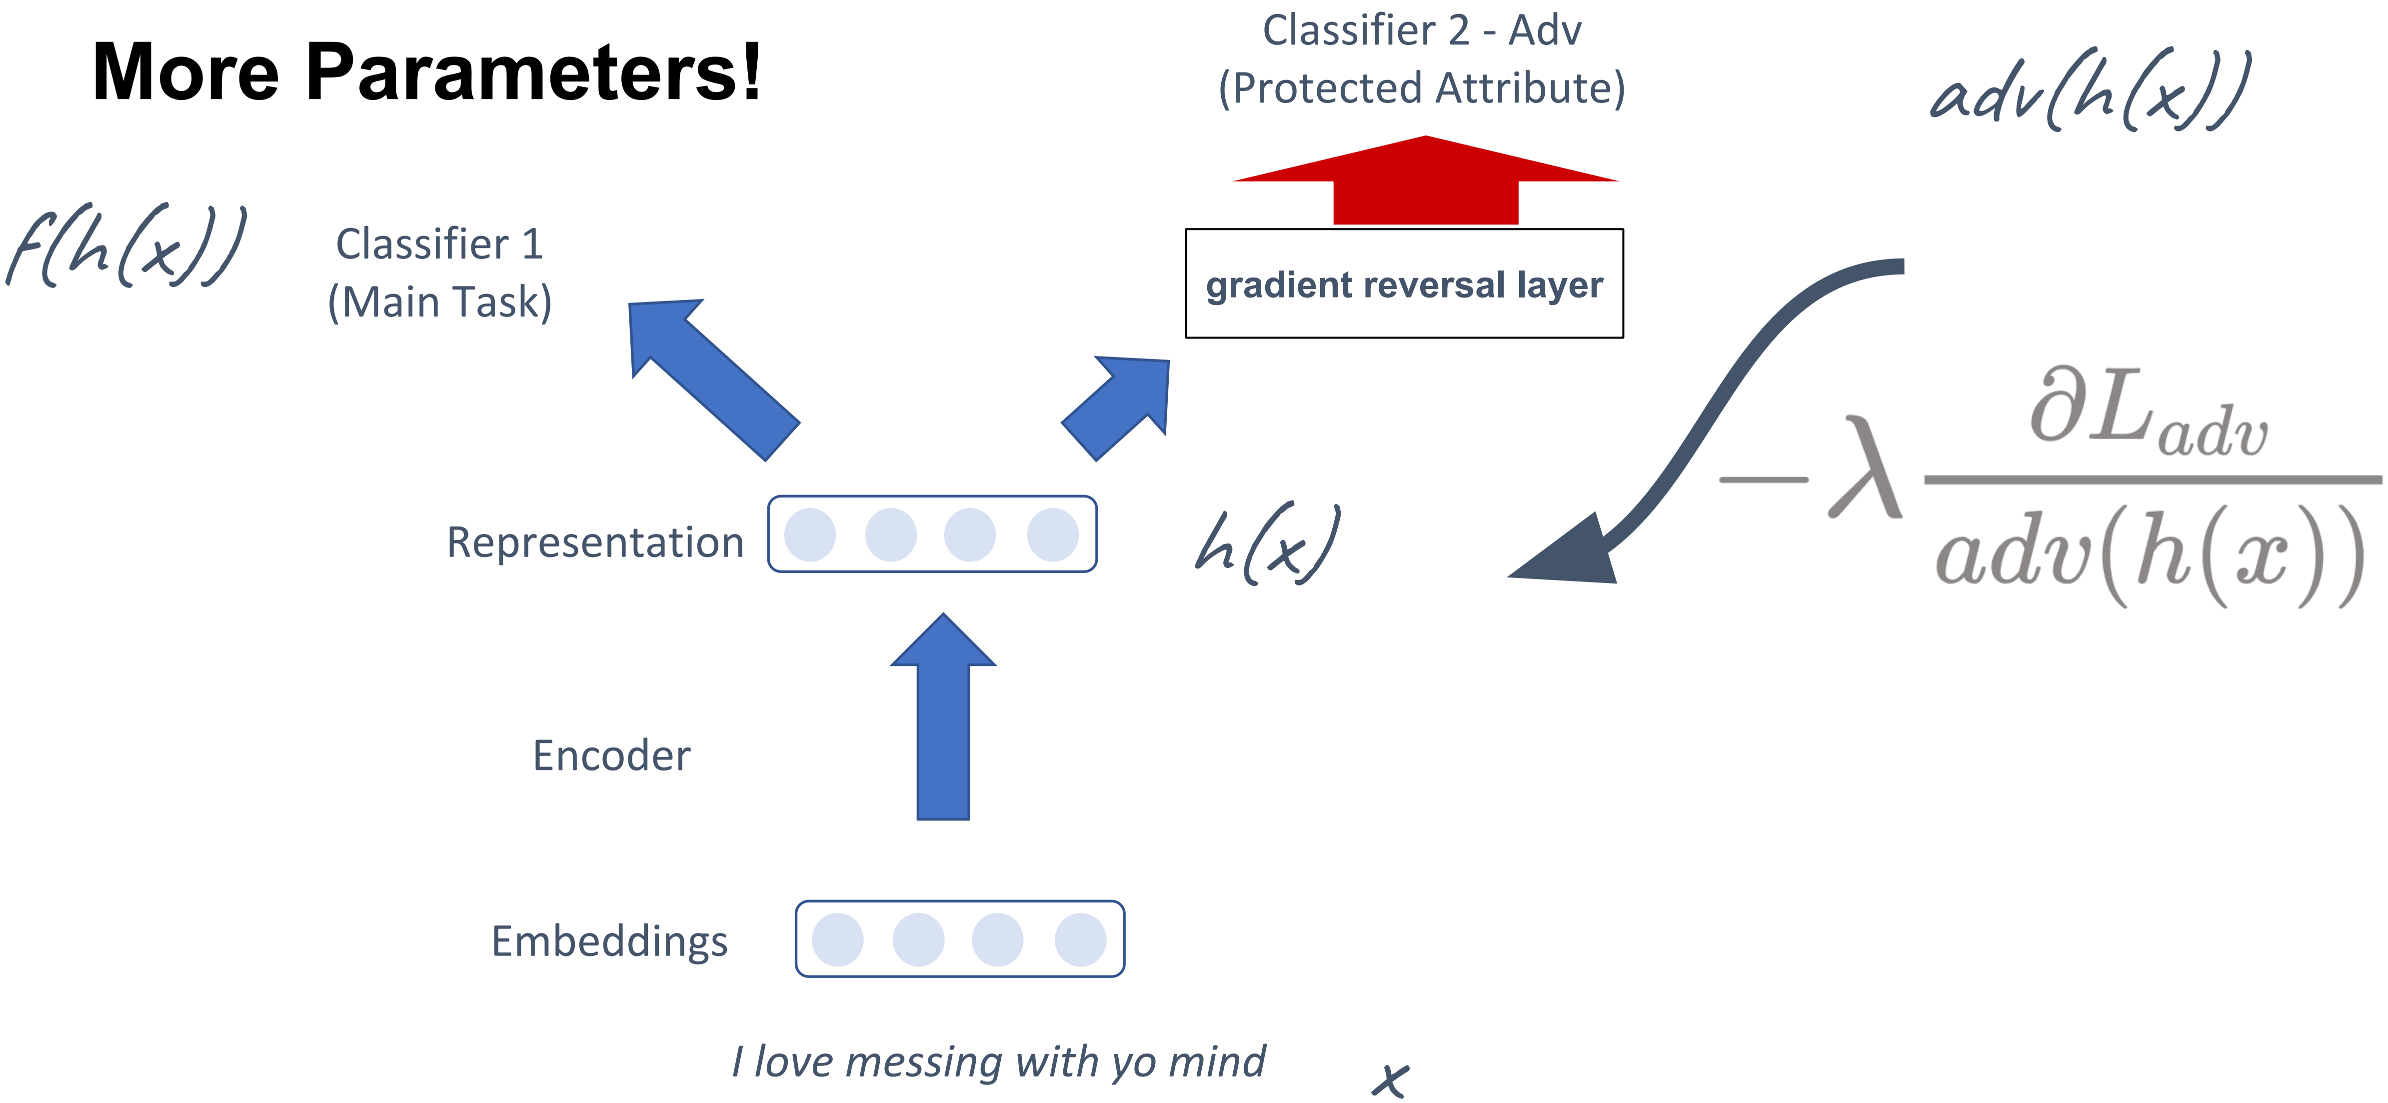

Stronger, Better, Bigger???

Baseline

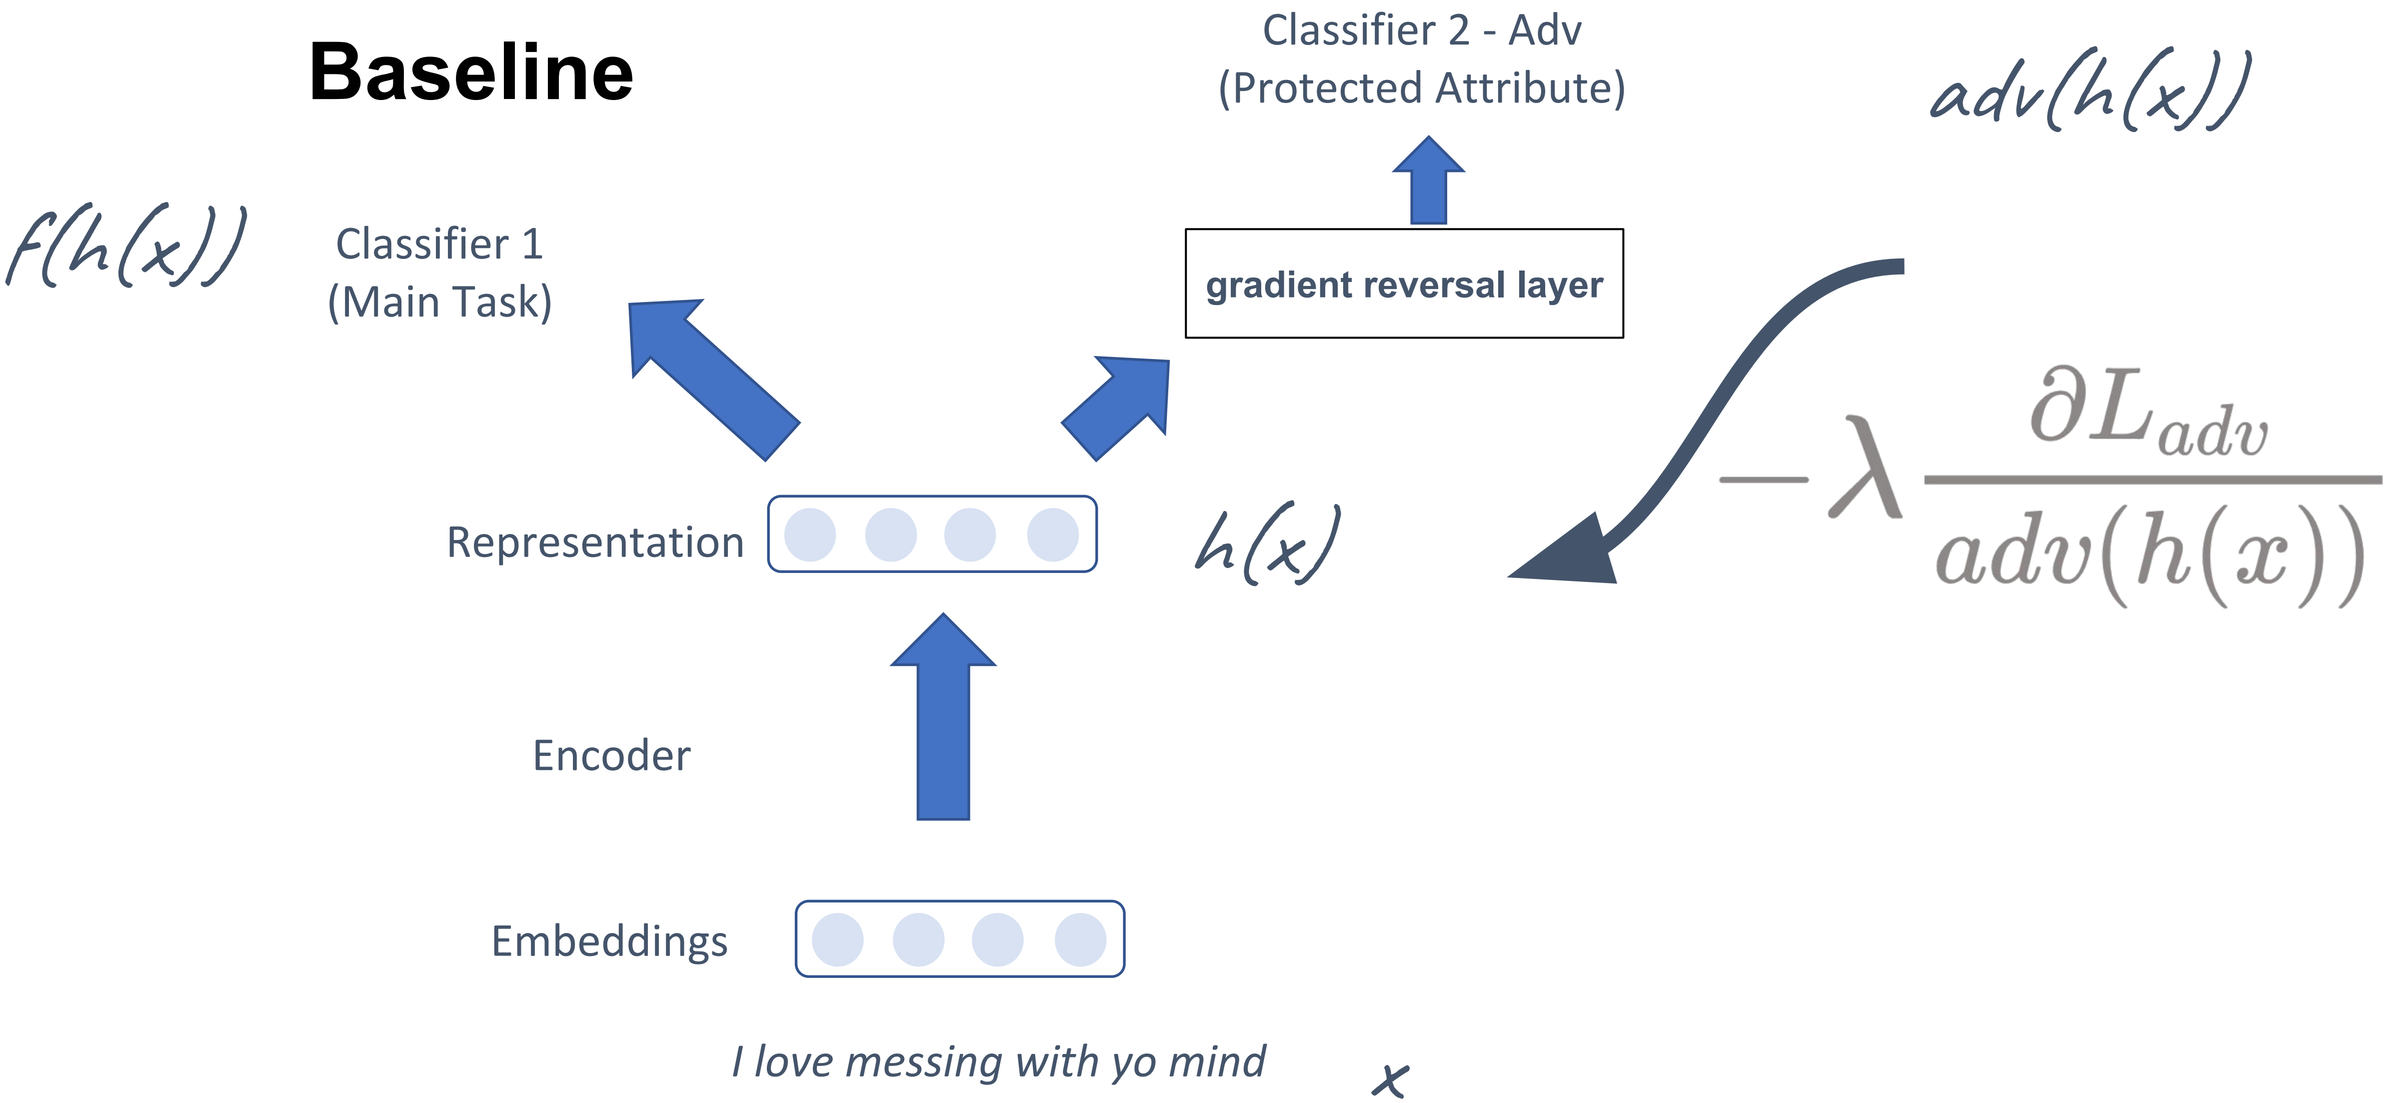

Stronger, Better, Bigger???

Bigger Weight!

Classifier 2 - Adv
(Protected Attribute)

$adv(h(x))$

gradient reversal layer

$f(h(x))$

Classifier 1
(Main Task)

$$-\lambda \frac{\partial L_{adv}}{adv(h(x))}$$

Representation

$h(x)$

Scale the reverse gradients

Encoder

Embeddings

I love messing with yo mind

x

Stronger, Better, Bigger???

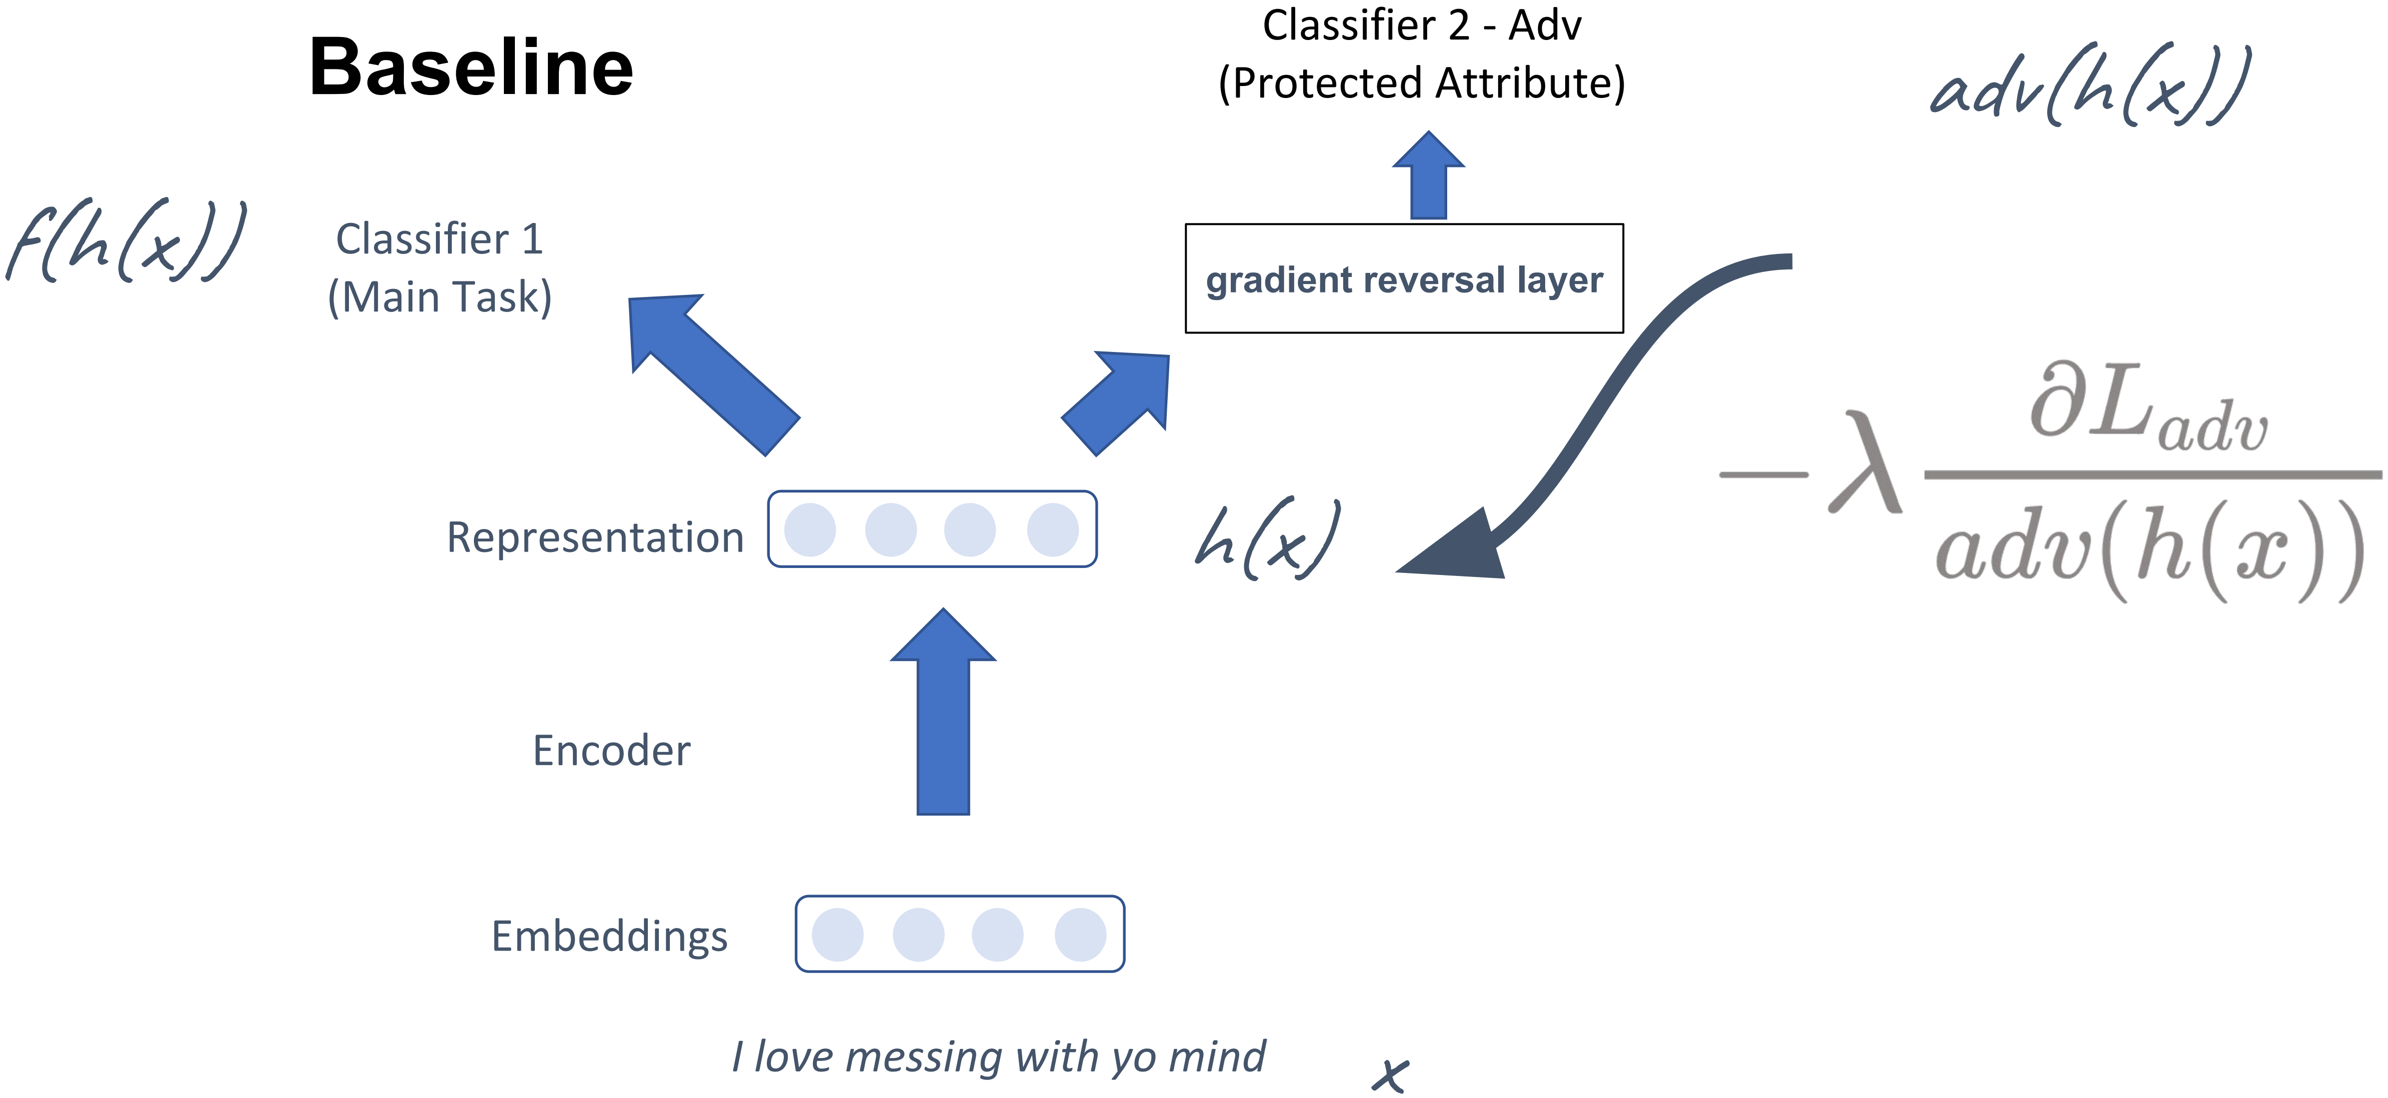

Stronger, Better, Bigger???

More Adversaries!

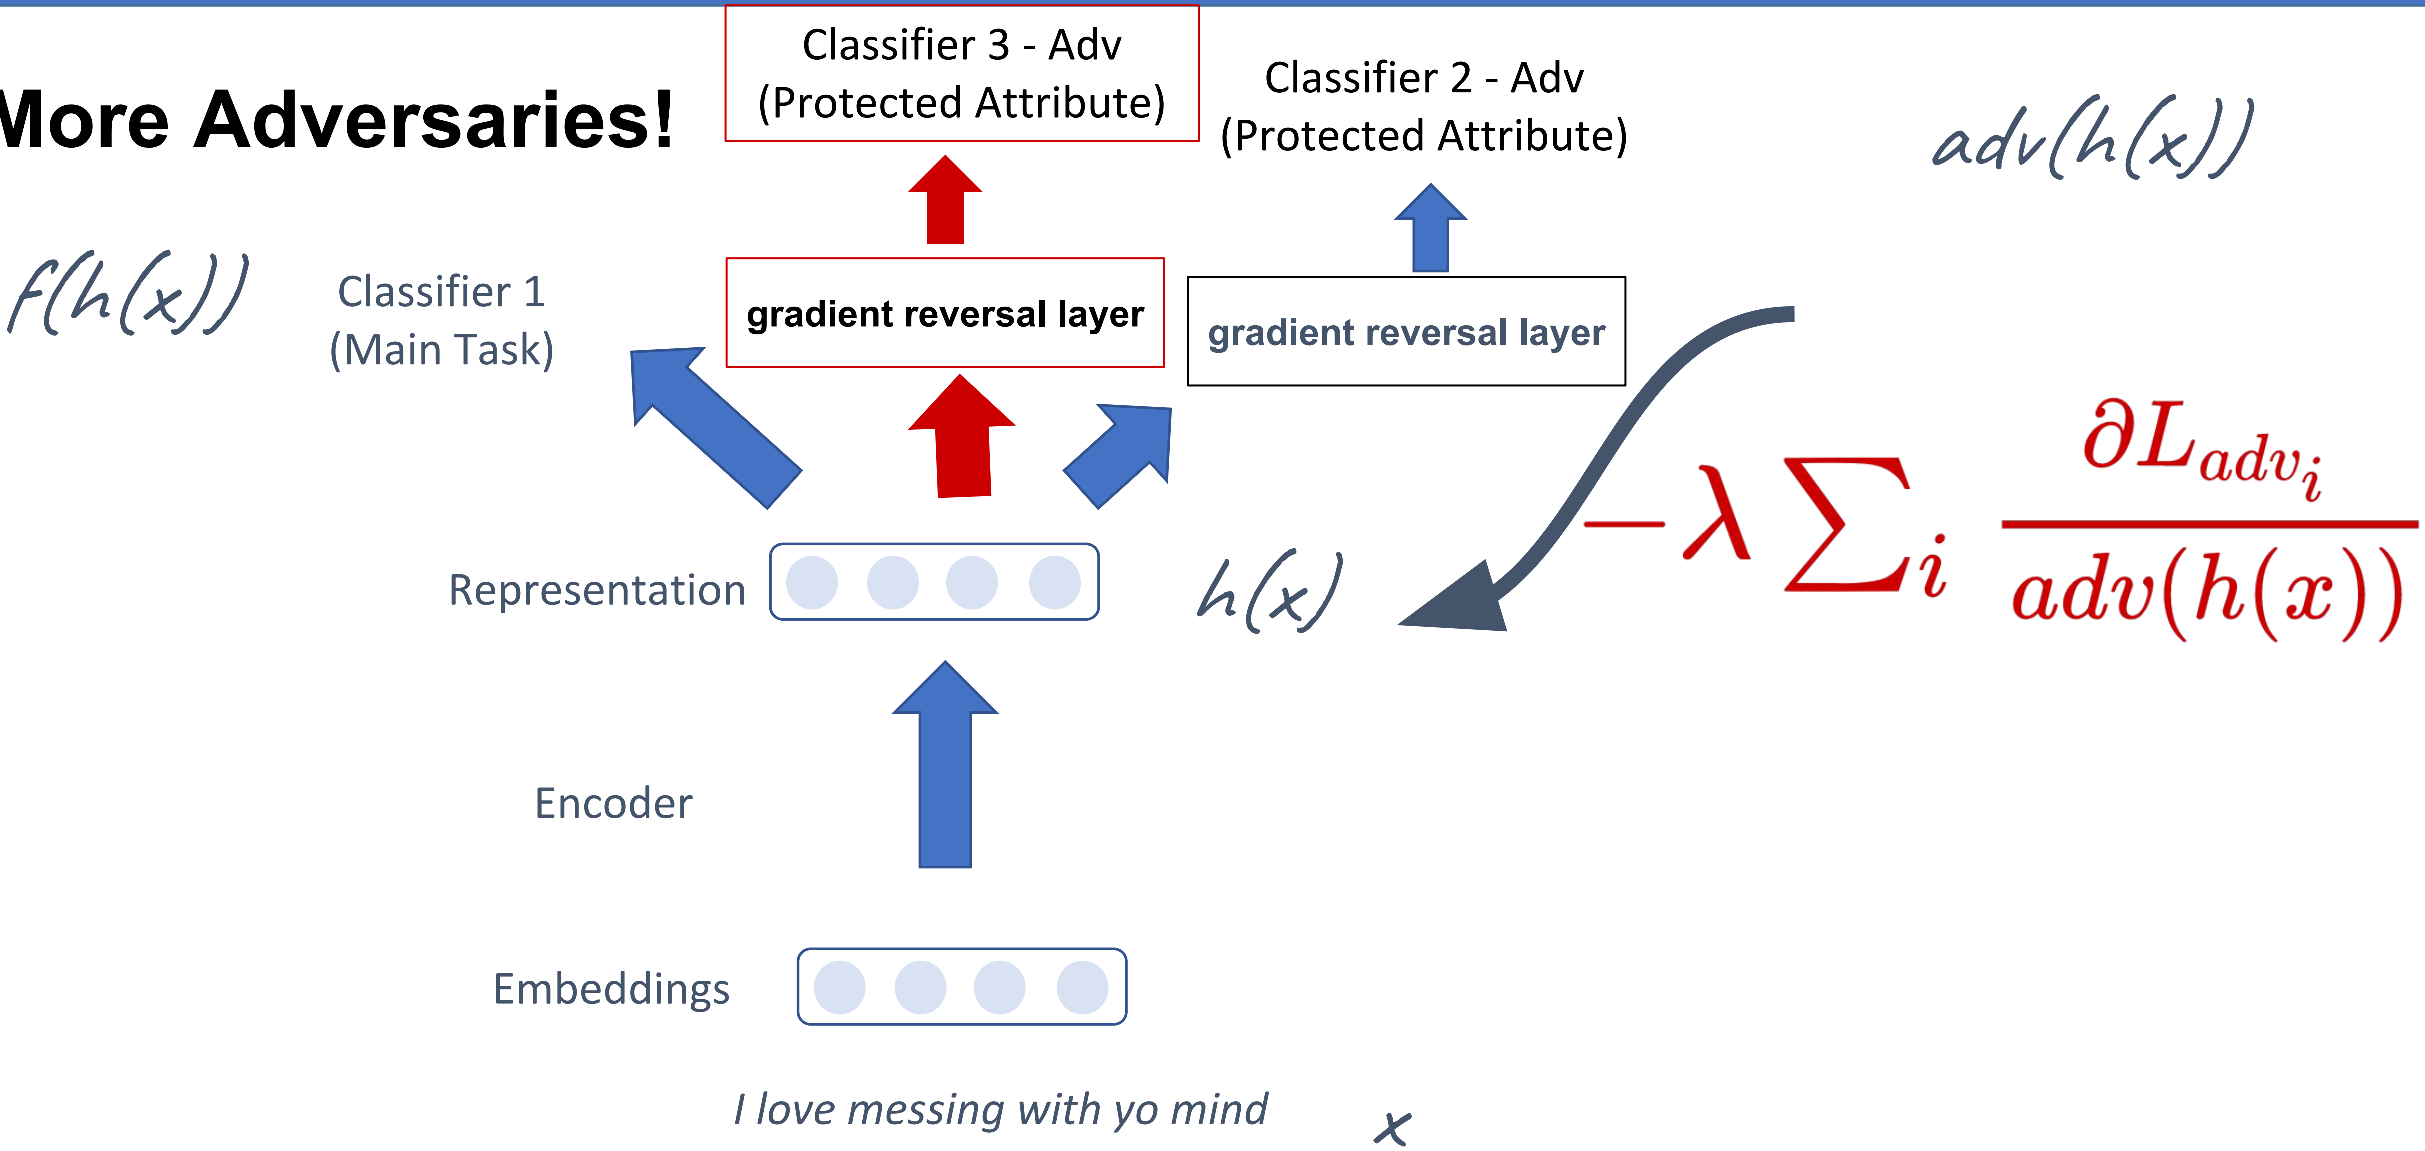

Stronger, Better, Bigger???

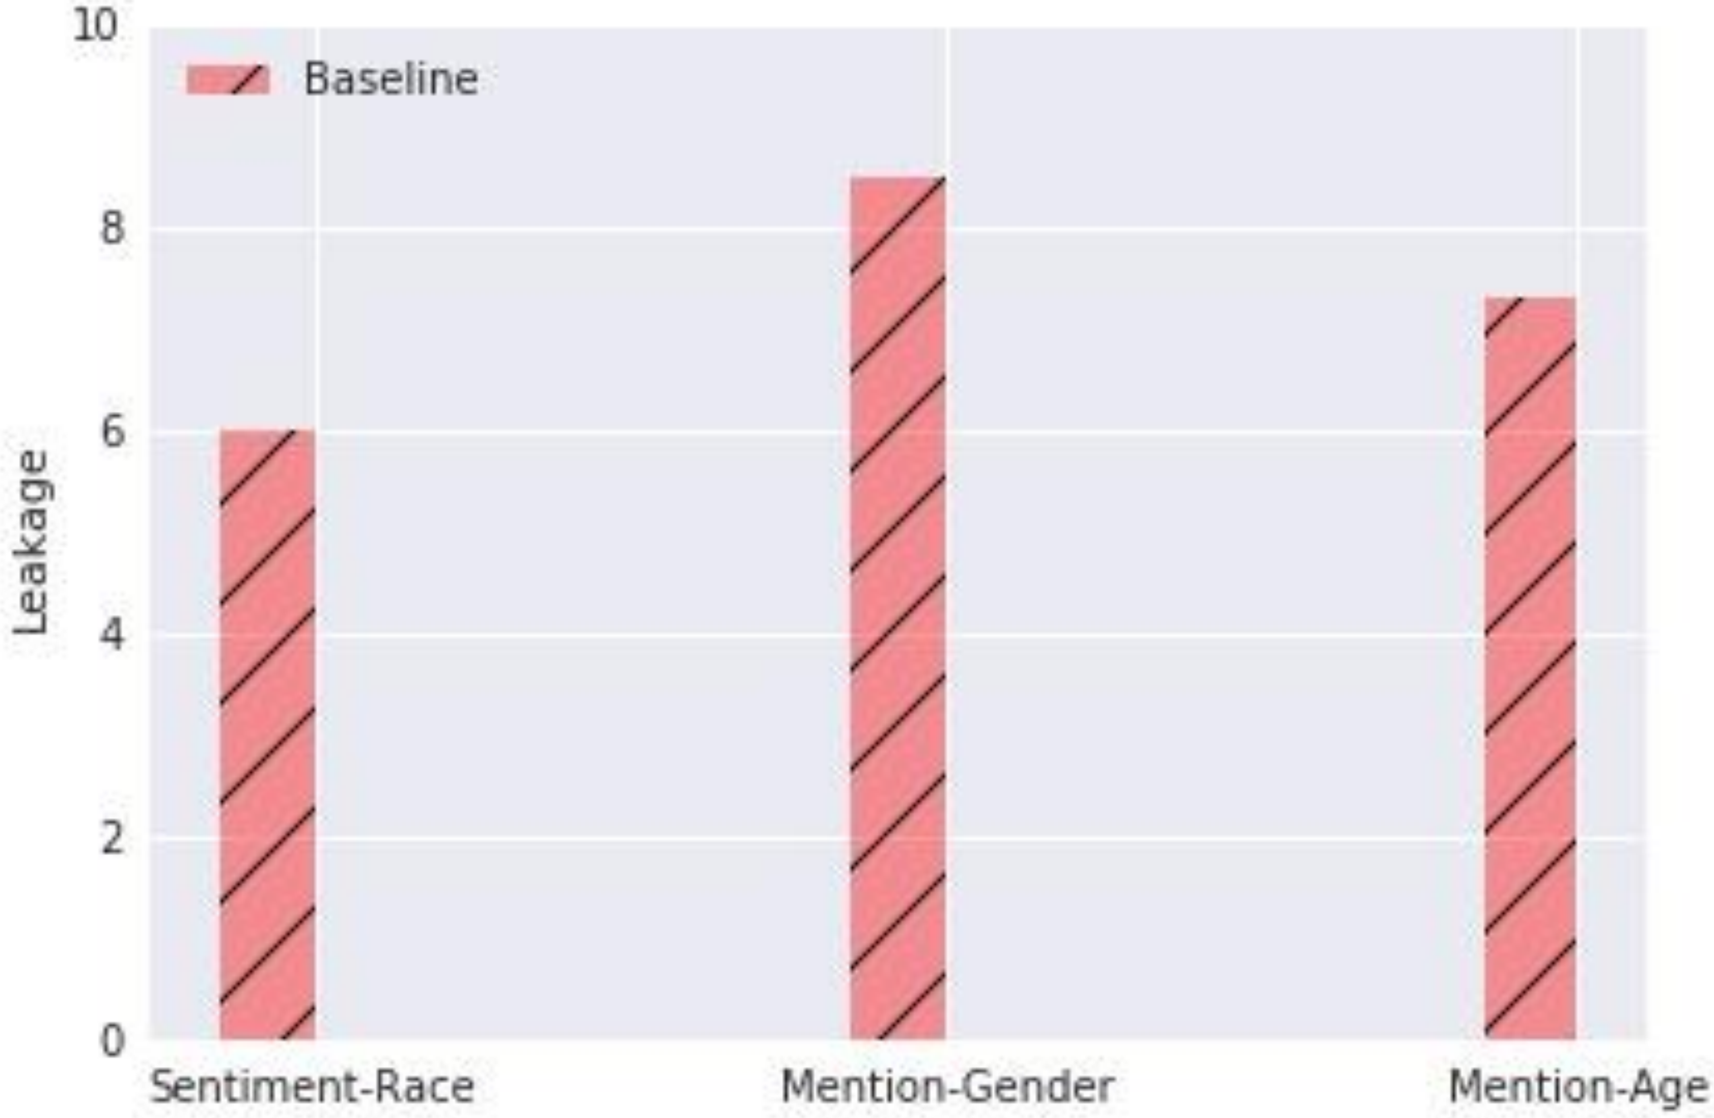

Stronger, Better, Bigger???

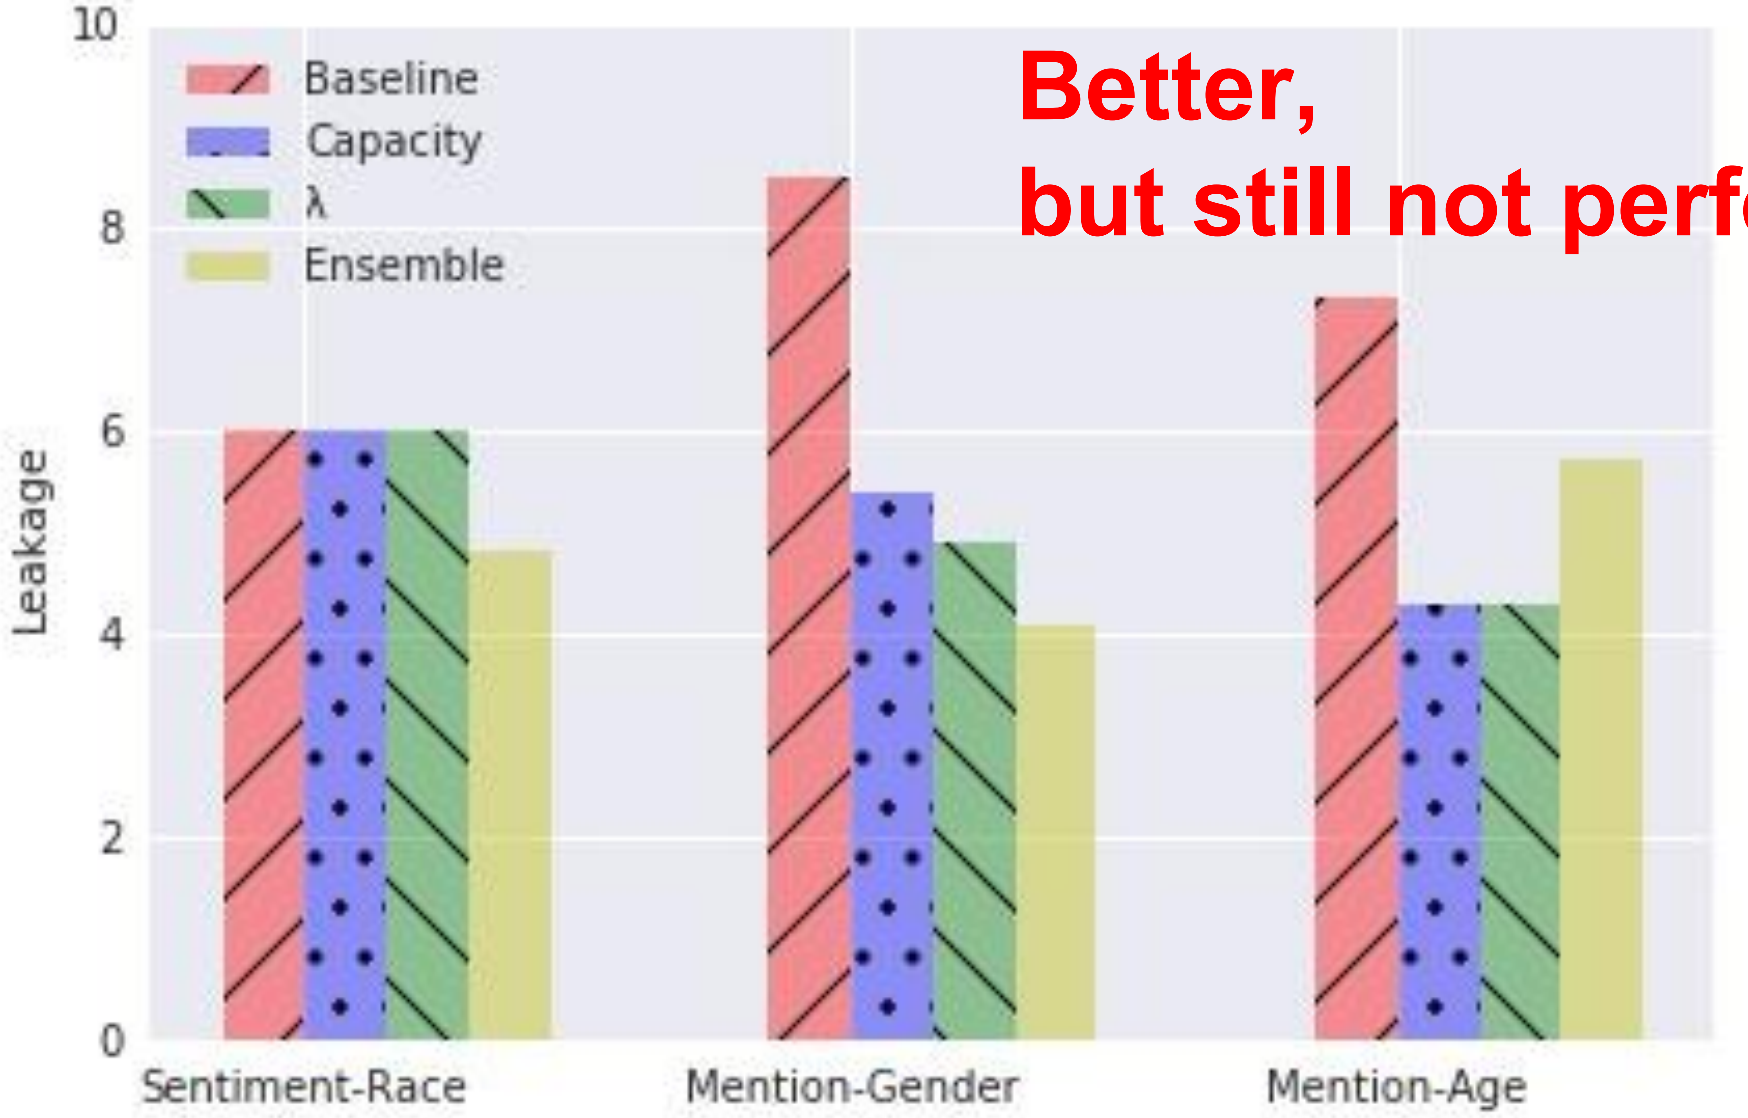

Better,
but still not perfect

Error Analysis

Persistent Examples

- What are the hard cases, which slip the adversary?
  - We trained the adversarial model 10 times (with random seeds)
  - then, trained the Attacker on each model
  - We collected all examples, which were consistently labeled correctly

Persistent Examples

- What are the hard cases, which slip the adversary?

Persistent Examples

AAE(“non-hispanic blacks”)
Enoy yall day
_ Naw im cool
My Brew Eatting
My momma Bestfrand died
Tonoght was cool

SAE (“non-hispanic whites”)
I want to be tan again
Why is it so hot in the house?!
I want to move to california
I wish I was still in Spain
Ahhhh so much homework.

More about the leakage origin can be found in the paper

Summary

- When training a text encoder for some task
  - Encoded vectors are also useful for predicting other things ("transfer learning")
  - Including things we did not want to encode ("leakage")
- **It is hard to completely prevent such leakage**
  - **Do not blindly trust adversarial training**
  - **Verify your model using an "Attacker"**

Thank you

Wait. I remember this thing called Overfitting

- We still have a problem
  - During training it seems that the information was removed
  - But the Attacker tells us another story
- Everything we reported was on the dev-set
- Is it possible that we just overfitted on the training-set?

Wait. I remember this thing called Overfitting

- “Adversary overfitting”:
  - Memorizing the training data
  - By removing all its sensitive information
  - While leaking in test time

Wait. I remember this thing called Overfitting

We trained on 90% on the “overfitted” training set, and tested the remaining 10%

Training Set

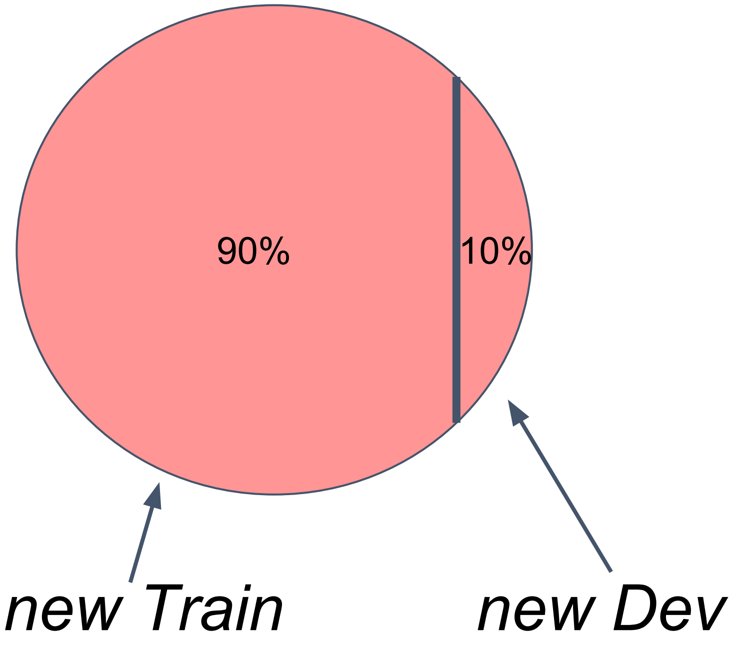

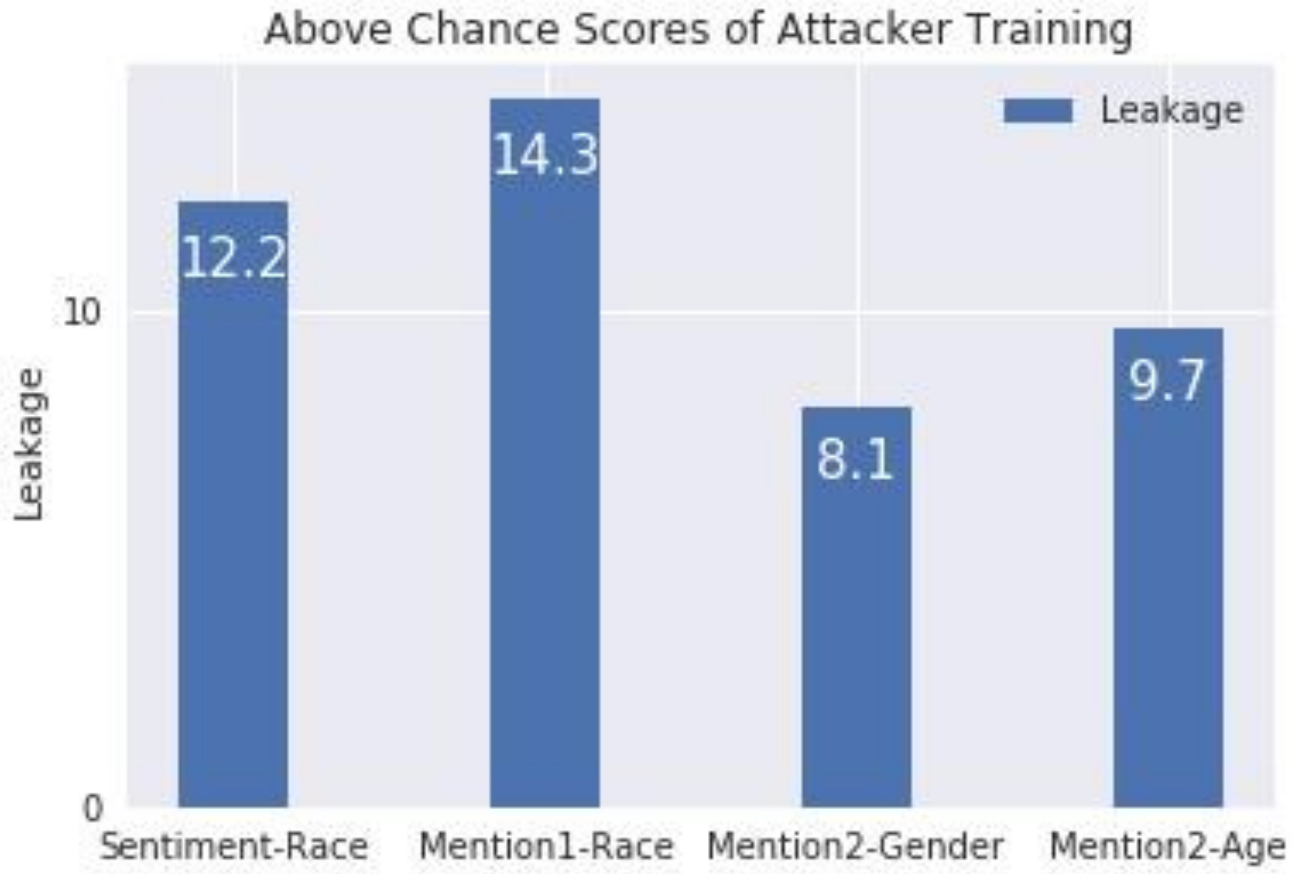

It is more than that

Few words about fairness

- Throughout this work, we aimed in achieving zero leakage, or in other words: *fairness by blindness*

- Many other definitions for "fairness" (>20)
- With 3 popular
  - *Demographic parity*
  - *Equality of Odds*
  - *Equality of Opportunity*

In the paper, we prove that in our setup (balanced data) these definitions are identical