CS 446: Machine Learning

Introduction to LBJava: a Learning Based Programming Language

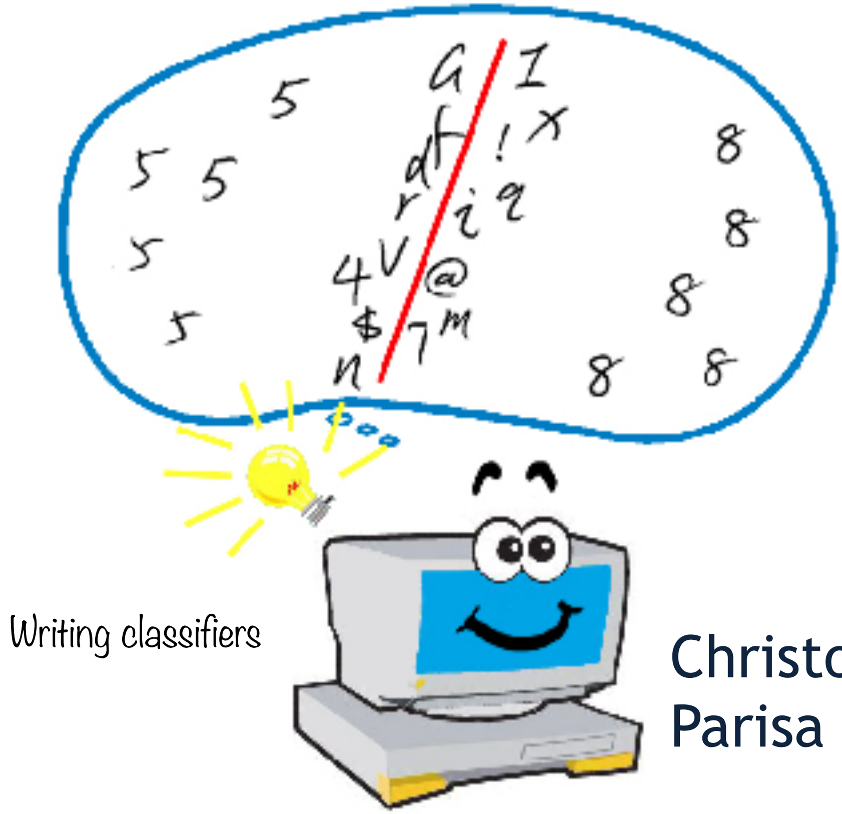

Christos Christodoulopoulos
Parisa Kordjamshidi

Motivation

Motivation

- You still have not learnt machine learning algorithms

Motivation

- You still have not learnt machine learning algorithms
  - But you can do cool things with the existing tools

Motivation

- You still have not learnt machine learning algorithms
  - But you can do cool things with the existing tools
  - And even earn money using it ;-)!!

Motivation

- You still have not learnt machine learning algorithms
  - But you can do cool things with the existing tools
  - And even earn money using it ;-)!!

Google DeepMind!

Yahoo Summly!

Tweeters WhetLab!

Motivation

- You still have not learnt machine learning algorithms
    - But you can do cool things with the existing tools
    - And even earn money using it ;-)!!

Google DeepMind!

Yahoo Summly!

Tweeters WhetLab!

How?

One interesting application:

Lets analyse tweets!

Data

Twitter posts

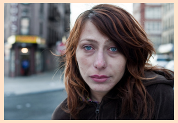

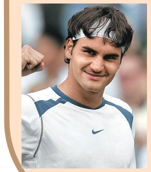

One interesting application:

Lets analyse tweets!

Data

Twitter posts

Heard someone sing a Christmas song, in the pub on Friday night. Give us a break!
Place: Illinois, USA, United States

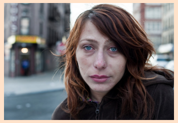

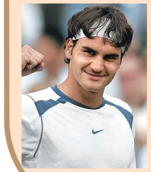

One interesting application:

Lets analyse tweets!

Data

Twitter posts

Heard someone sing a Christmas song, in the pub on Friday night. Give
Place: Illinois, USA, United States

happy tweet

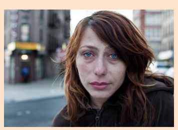

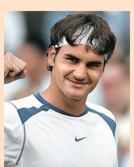

One interesting application:

Lets analyse tweets!

Data

Twitter posts

Heard someone sing a Christmas song, in the pub on Friday night. Give
Place: Illinois, USA, United States

happy tweet

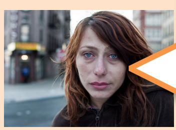

My mom just dragged me to Walgreens and forced me to get a flu shot and then she told me it was just like mother–daughter tattoos #help
Place: Illinois, USA, United States

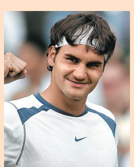

One interesting application:

Lets analyse tweets!

Data

Twitter posts

Heard someone sing a Christmas song, in the pub on Friday night. Give
Place: Illinois, USA, United States

happy tweet

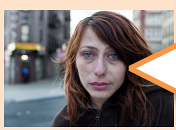

My mom just dragged me to Walgreens and forced me to get a f
she told me it was just like mother–daughter tattoos #help
Place: Illinois, USA, United States

unhappy tweet

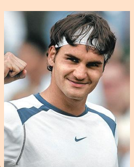

One interesting application:

Lets analyse tweets!

Data

Twitter posts

Heard someone sing a Christmas song, in the pub on Friday night. Give a
Place: Illinois, USA, United States

happy tweet

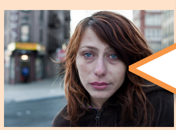

My mom just dragged me to Walgreens and forced me to get a fl
she told me it was just like mother-daughter tattoos #help
Place: Illinois, USA, United States

unhappy tweet

just wanna leave these past in the past and move on
Place: Edinburgh, Scotland, United Kingdom

One interesting application:

Lets analyse tweets!

Data

Twitter posts

Heard someone sing a Christmas song, in the pub on Friday night. Give
Place: Illinois, USA, United States

happy tweet

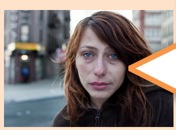

My mom just dragged me to Walgreens and forced me to get a f
she told me it was just like mother-daughter tattoos #help
Place: Illinois, USA, United States

unhappy tweet

just wanna leave these past in the past and move on
Place: Edinburgh, Scotland, United Kingdom

unhappy tweet

One interesting application:

Lets analyse tweets!

Meta analysis: which location is happier?

Twitter posts

Heard someone sing a Christmas song, in the pub on Friday night. Give
Place: Illinois, USA, United States

happy tweet

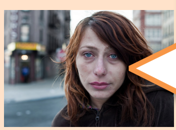

My mom just dragged me to Walgreens and forced me to get a f
she told me it was just like mother–daughter tattoos #help
Place: Illinois, USA, United States

unhappy tweet

just wanna leave these past in the past and move on
Place: Edinburgh, Scotland, United Kingdom

unhappy tweet

Our application
Sentiment analysis of tweets!

What are the steps?

Our application
Sentiment analysis of tweets!

What are the steps?

- Create examples: Let's get people label some tweets with positive/negative labels.

Our application

Sentiment analysis of tweets!

What are the steps?

- Create examples: Let's get people label some tweets with positive/negative labels.
- Look at examples: Find out the influencing features in realizing the sense of a tweet.

Our application

Sentiment analysis of tweets!

What are the steps?

- Create examples: Let's get people label some tweets with positive/negative labels.
- Look at examples: Find out the influencing features in realizing the sense of a tweet.
- Choose a learning technique and train: Apply the technique on the labeled data and build a model.

Our application

Sentiment analysis of tweets!

What are the steps?

- Create examples: Let's get people label some tweets with positive/negative labels.
- Look at examples: Find out the influencing features in realizing the sense of a tweet.
- Choose a learning technique and train: Apply the technique on the labeled data and build a model.
- Use the model: Make predictions on the unlabeled tweets later.

Our application
Sentiment analysis of tweets!

What are the steps?

- Create examples: Let’s get people label some tweets with positive/negative labels.
- Look at examples: Find out the influencing features in realizing the sense of a tweet.
- Choose a learning technique and train: Apply the technique on the labeled data and build a model.
- Use the model: Make predictions on the unlabeled tweets later.
- !BTW: Test the trained model using a subset of your labeled data, don't use that subset for training. You need to see how well your model can work.

Our application
Sentiment analysis of tweets!

What are the steps?

- **Create examples**: Let's get people label some tweets with positive/ negative labels.
- **Look at examples**: Find out the influencing features in realizing the sense of a tweet.
- **Learning technique** and train: Apply the technique on the labeled data and build a model.
- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** the trained model using a subset of your labeled data, don't use that for training. You need to see how well your model can work.

t analysis of tweets!

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

- **Create examples**: Let’s get people label some tweets with positive/ negative labels.
- **Look at examples**: Find out the influencing features in realizing the sense of a tweet.
- **Learning technique** and train: Apply the technique on the labeled data and build a model.
- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** the trained model using a subset of your labeled data, don't use that for training. You need to see how well your model can work.

t analysis of tweets!

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

What are the steps:

- **Create examples**: Let's
  negative labels.
- **Look at examples**: Fi
  sense of a tweet.

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases.

- **Learning technique** and train: Apply the technique on the labeled data and build a model.
- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** the trained model using a subset of your labeled data, don't use that for training. You need to see how well your model can work.

…t analysis of tweets!

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

- **Create examples**: Let's … negative labels.
- **Look at examples**: Fi… sense of a tweet.

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases.

- **Learning technique** and trai… …beled data and build a model.

SVMs? Naive Bayes?
Decision trees?
Sparse Perceptrons? …

- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** the trained model using a subset of your labeled data, don't use that for training. You need to see how well your model can work.

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

t analysis of tweets!

- **Create examples**: Let's

  negative labels.

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases.

- **Look at examples**: Fi

  sense of a tweet.

- **Learning technique** and trai beled data and build a model.

SVMs? Naive Bayes?
Decision trees?
Sparse Perceptrons? …

- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** t

  use that for t

Future tweets also should be represented with same representation (features) as training.

... t analysis of tweets!

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

What are the steps?

- **Create examples**: Let's ... negative labels.
- **Look at examples**: Fi... sense of a tweet.
- **Learning technique** and trai... beled data and build a model.
- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** t...

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases.

SVMs? Naive Bayes?
Decision trees?
Sparse Perceptrons? ...

Future tweets also should be represented with same representation (features) as training.

use the commonly used evaluation metrics.

…t analysis of tweets!

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

1

- **Create examples**: Let's … negative labels.
- **Look at examples**: Fi… sense of a tweet.
- **Learning technique** and trai… …beled data and build a model.
- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** t…

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases.

SVMs? Naive Bayes?
Decision trees?
Sparse Perceptrons? …

Future tweets also should be represented with same representation (features) as training.

use the commonly used evaluation metrics.

t analysis of tweets!

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

1

- **Create examples**: Let's negative labels.
- **Look at examples**: Fi sense of a tweet.
- **Learning technique** and trai beled data and build a model.
- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** t

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases.

2

SVMs? Naive Bayes?
Decision trees?
Sparse Perceptrons? …

Future tweets also should be represented with same representation (features) as training.

use the commonly used evaluation metrics.

...t analysis of tweets!

What are the steps?

- **Create examples**: Let's ... negative labels.
- **Look at examples**: Fi... sense of a tweet.
- **Learning technique** and trai... ...beled data and build a model.
- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** t...

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

1

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases.

2

SVMs? Naive Bayes? Decision trees? Sparse Perceptrons? …

3

Future tweets also should be represented with same representation (features) as training.

use the commonly used evaluation metrics.

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good!

1

t analysis of tweets!

- **Create examples**: Let's negative labels.

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases.

2

- **Look at examples**: F sense of a tweet.

- **Learning technique** and trai beled data and build a model.

SVMs? Naive Bayes?
Decision trees?
Sparse Perceptrons? ...

3

- **Use the model**: Make predictions on the unlabeled tweets later.

- !BTW: **Test** t

Future tweets also should be represented with same representation (features) as training.

5

use the commonly used evaluation metrics.

t analysis of tweets!

There is already a large number of tweets annotated by human and the data is publicly available. this step is done! Good! 1

- **Create examples**: Let's ... negative labels.
- **Look at examples**: ... sense of a tweet.

Represent each tweet simply with contained words. We can learn from these words occurrences, for example cry, sad, horrible, nice, pleasant, etc can tell us the sense of the sentence in many cases. 2

- **Learning technique** and trai... ...beled data and build a model.

SVMs? Naive Bayes?
Decision trees?
Sparse Perceptrons? … 3

- **Use the model**: Make predictions on the unlabeled tweets later.
- !BTW: **Test** t...

Future tweets also should be represented with same representation (features) as training. 5

use the commonly used evaluation metrics. 4

Our application
Sentiment analysis of tweets!

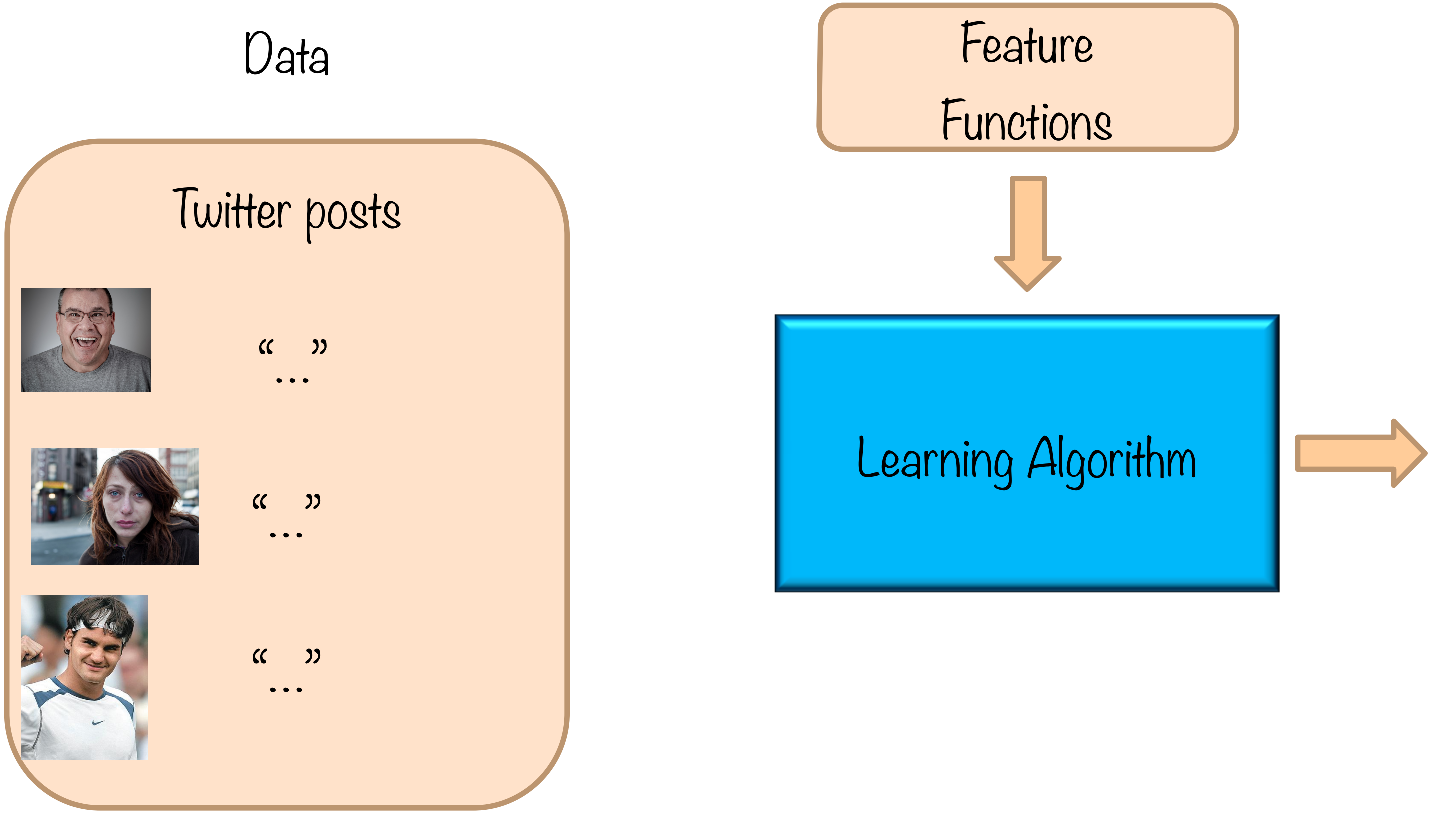

Our application

Sentiment analysis of tweets!

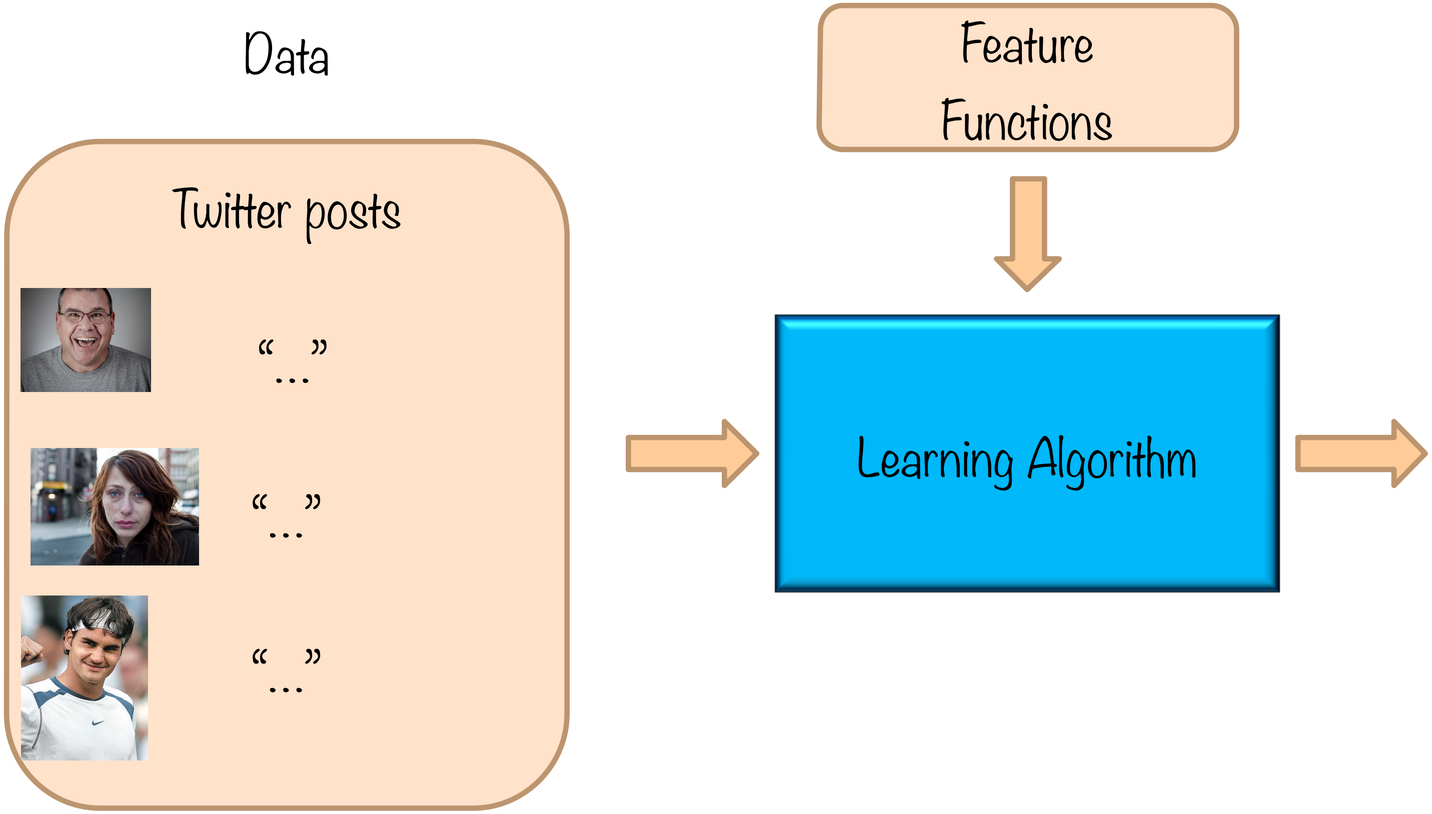

Our application
Sentiment analysis of tweets!

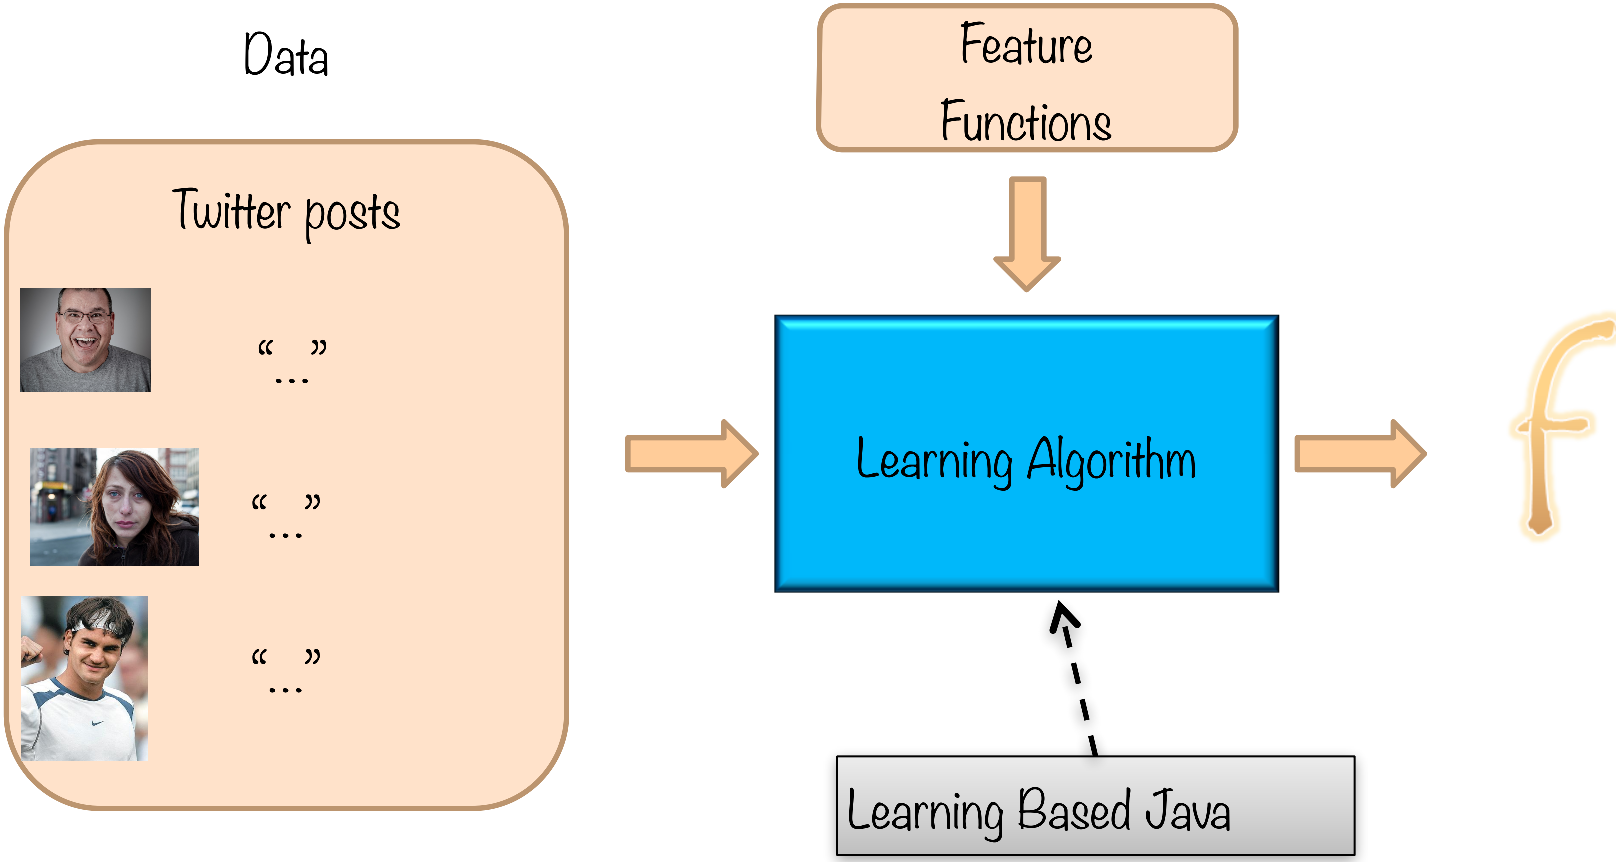

Our application

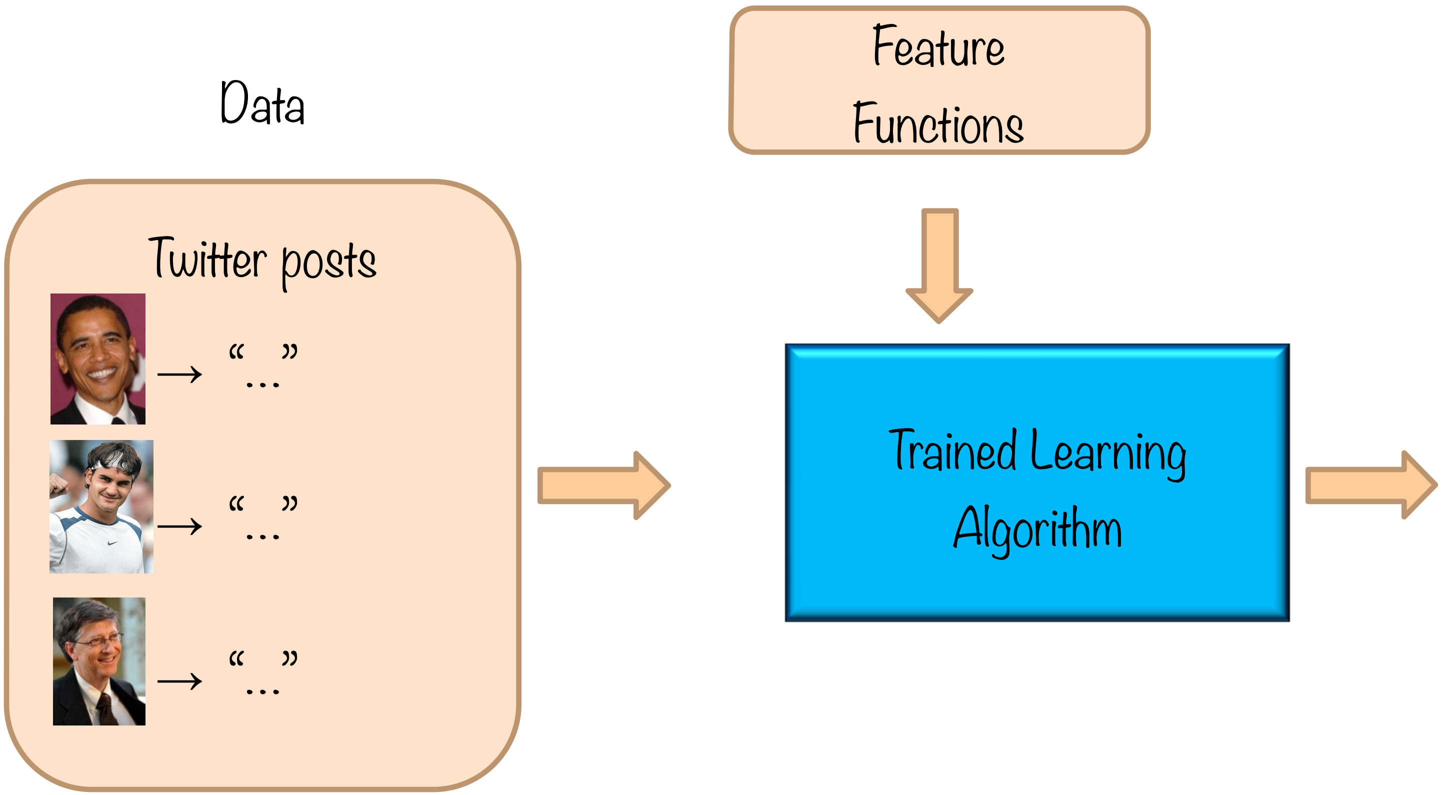

Our application

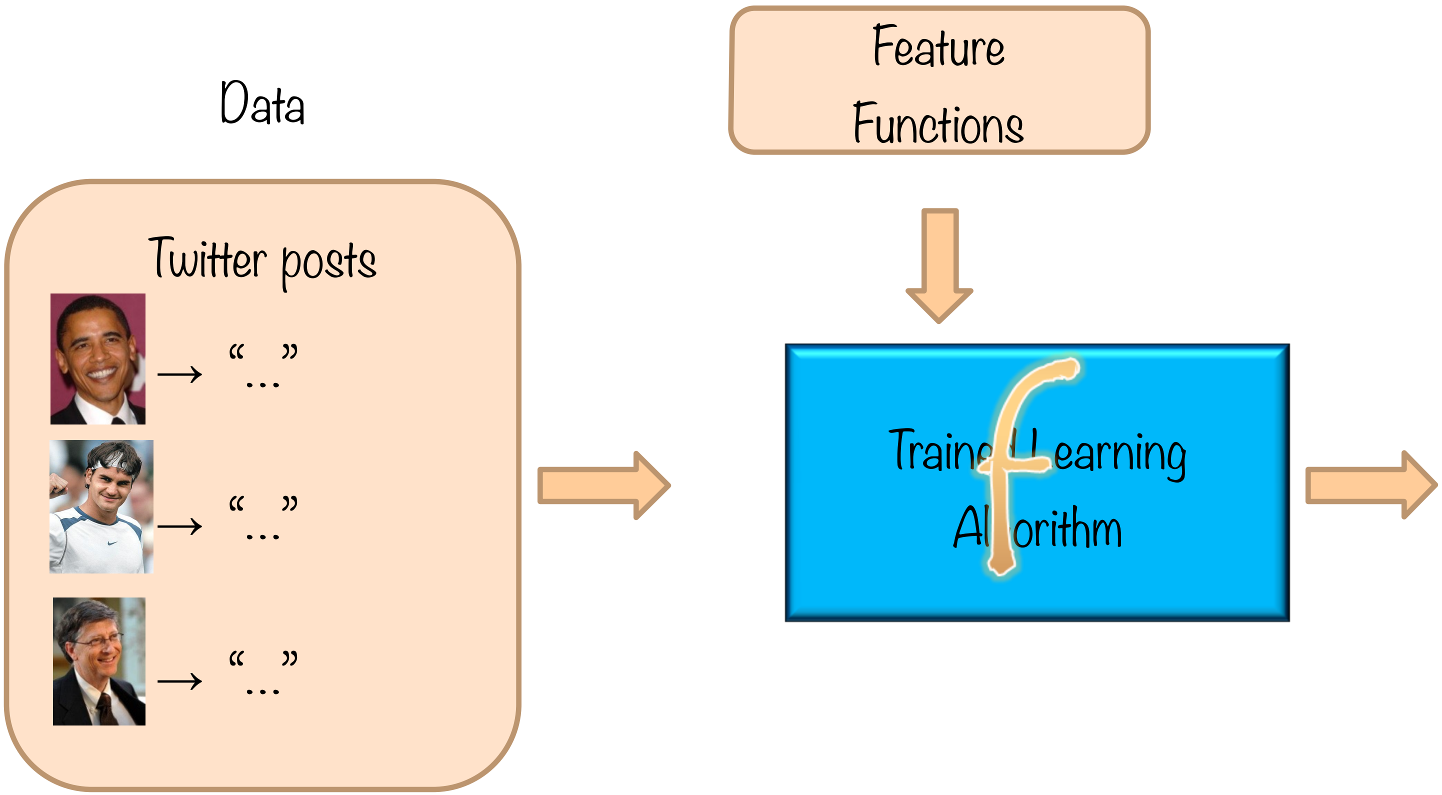

Our application

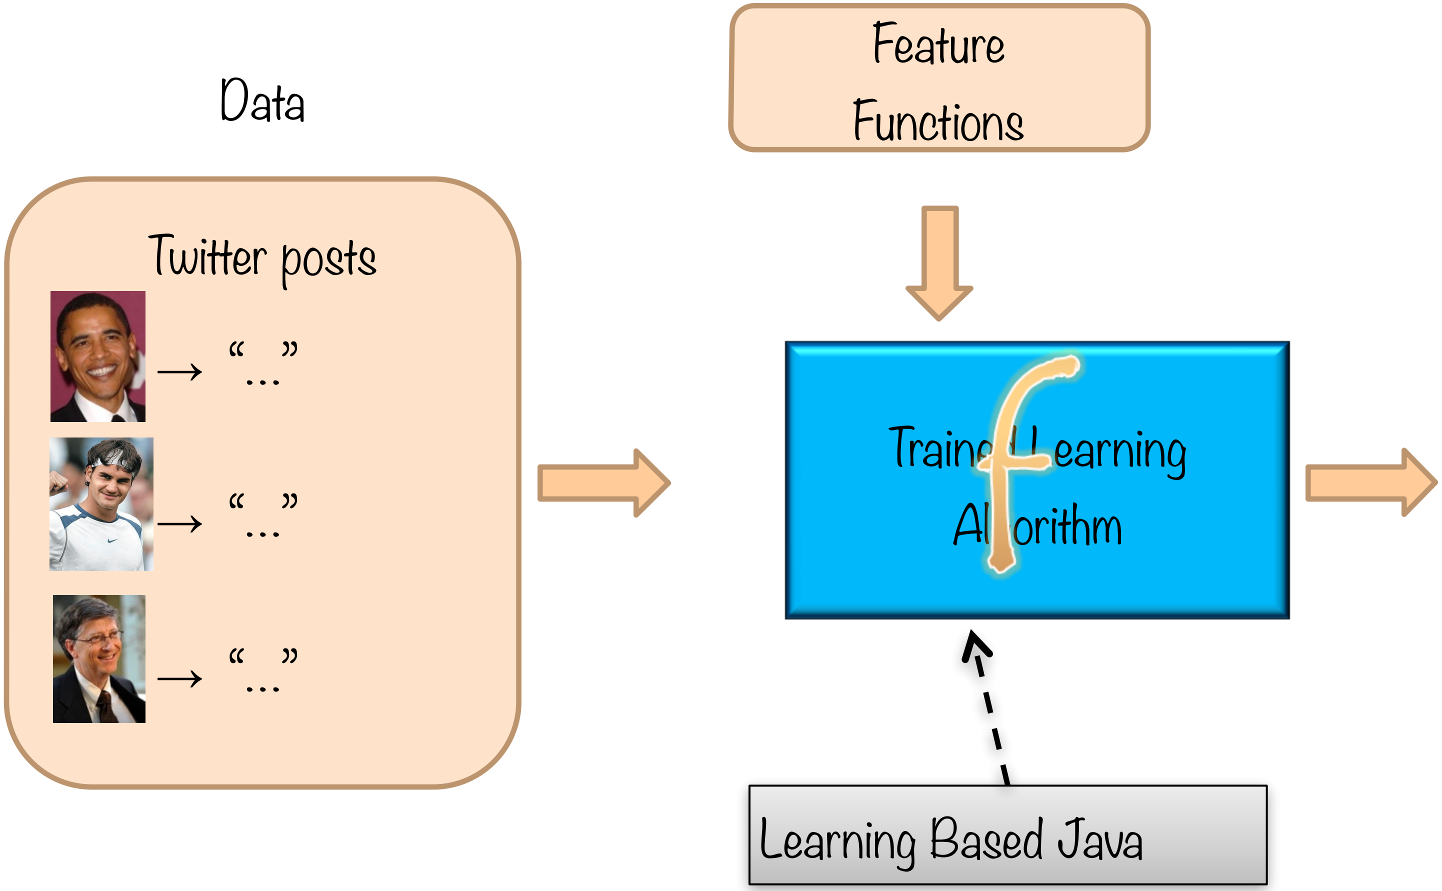

Our application

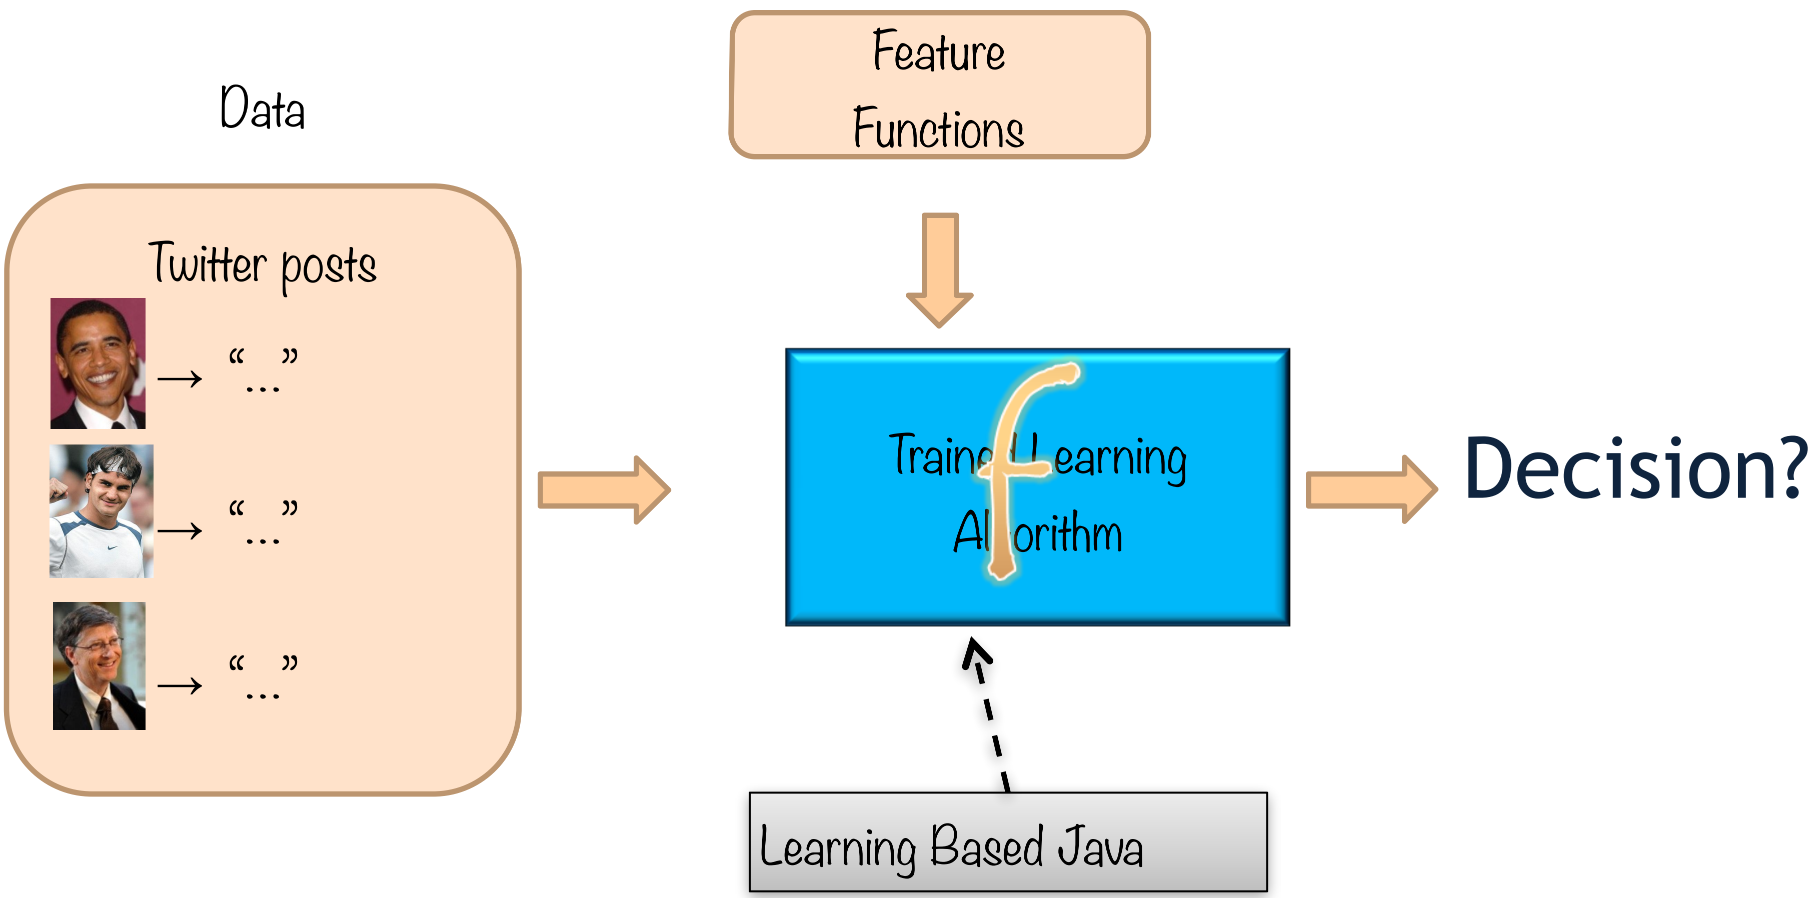

Our application

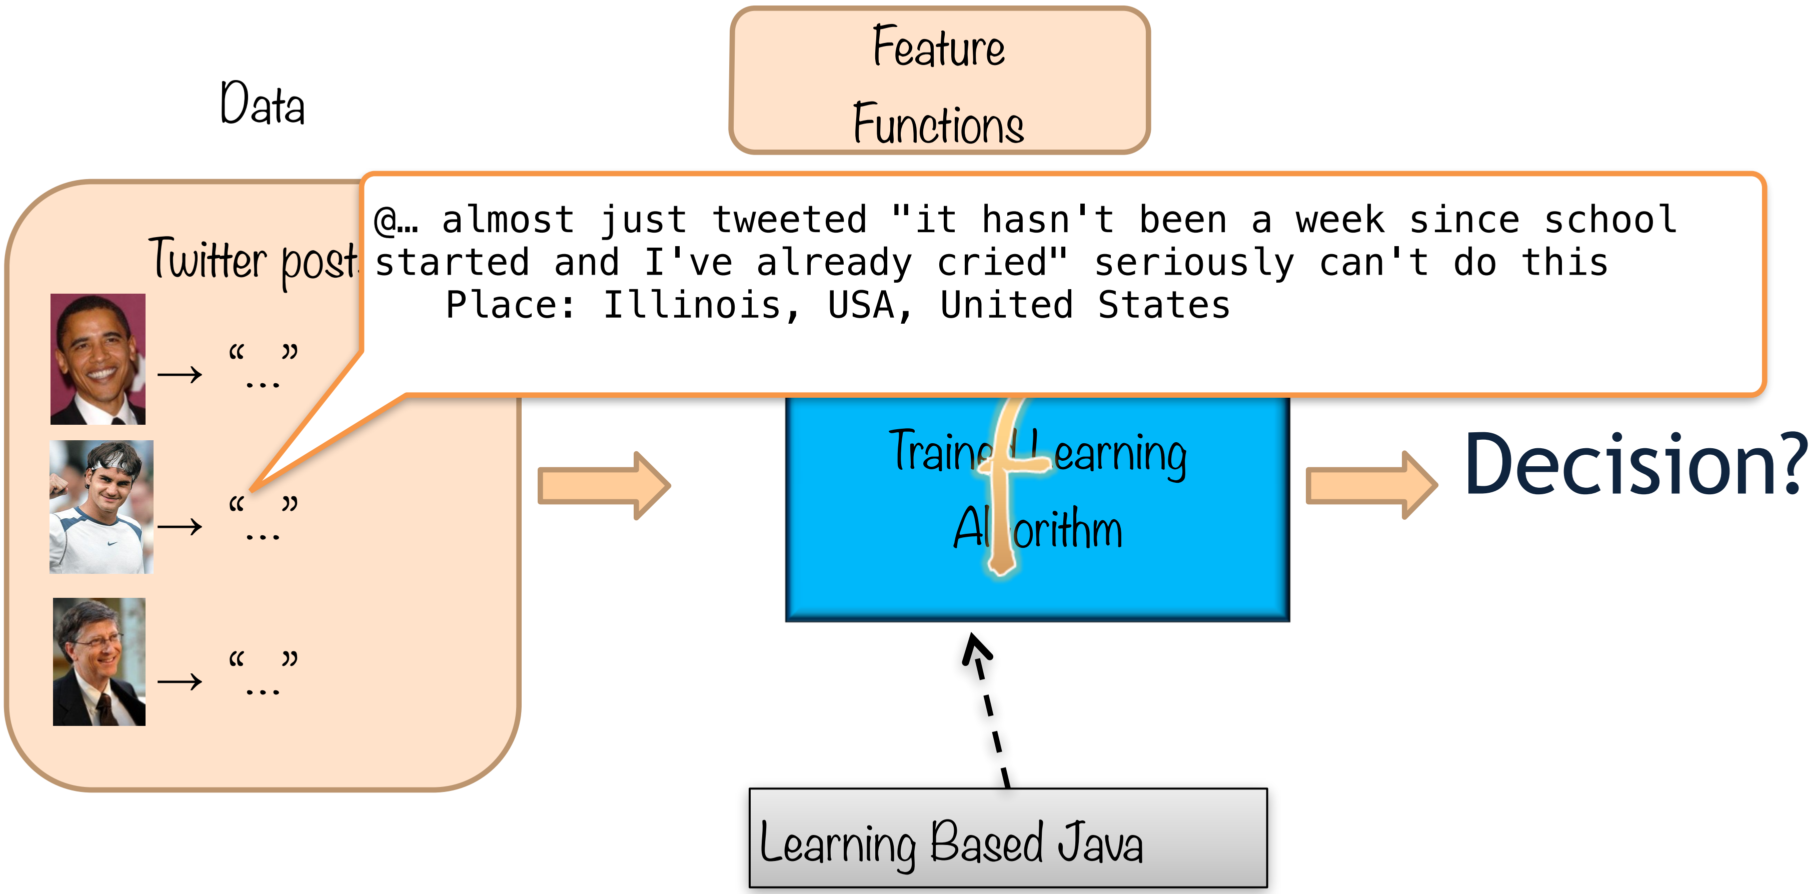

Our application

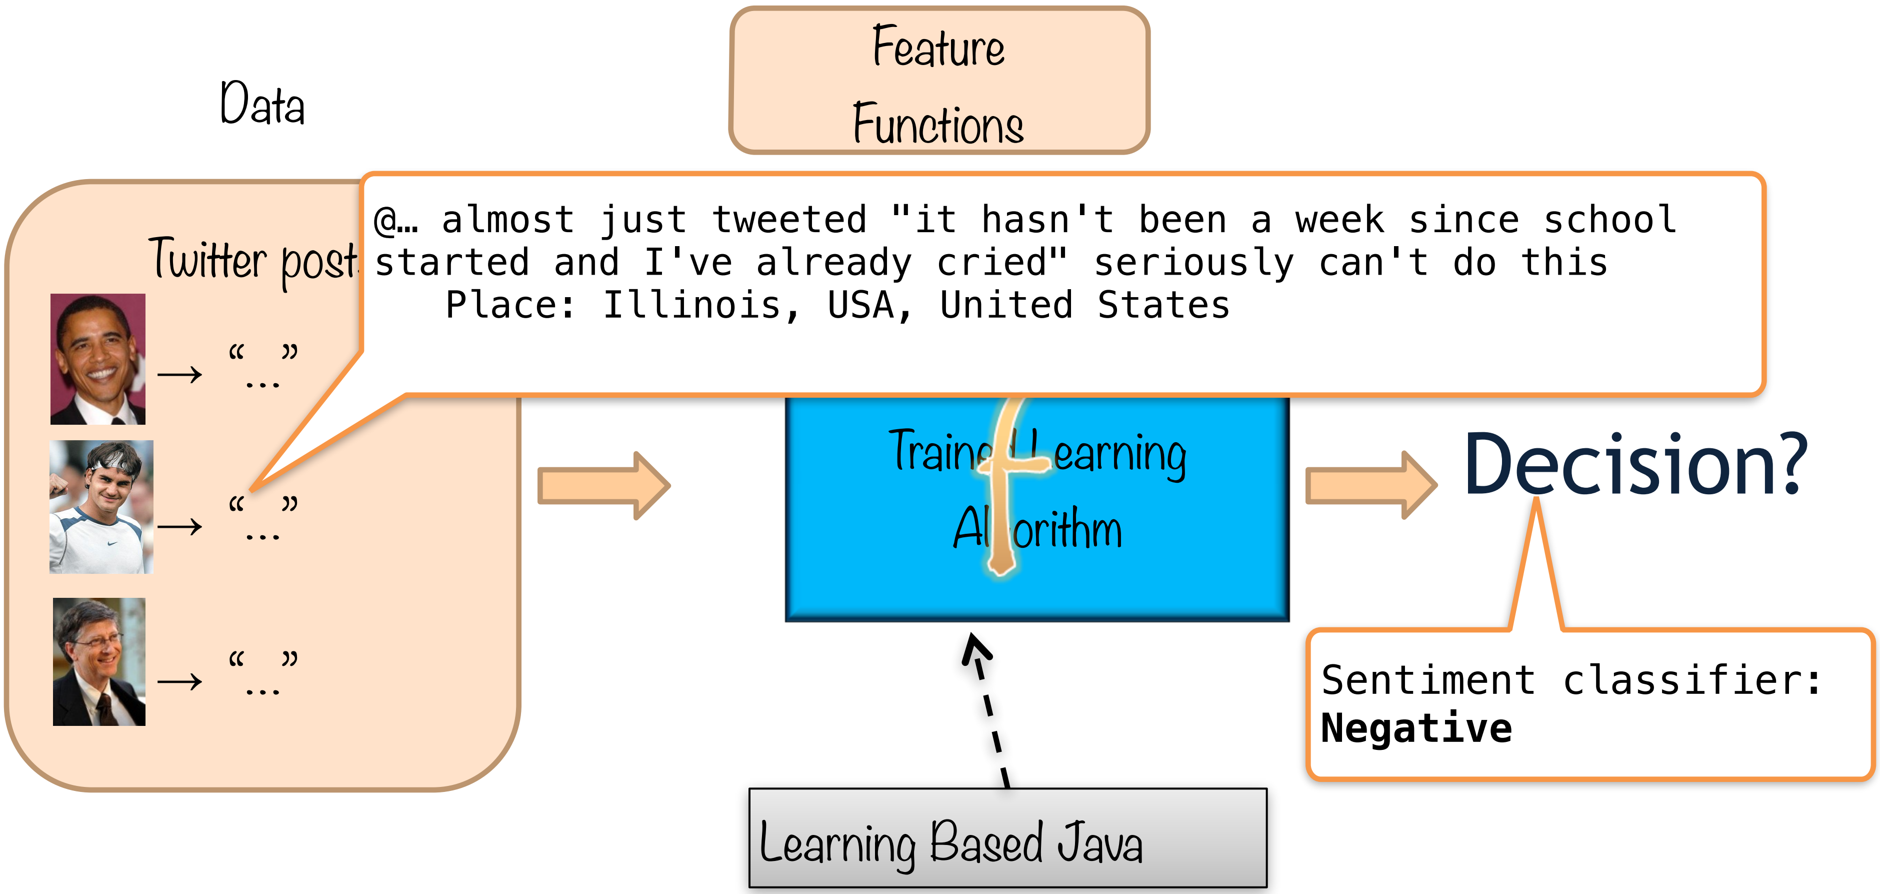

What is Learning Based Java?

What is Learning Based Java?

- A modeling language for learning and inference

What is Learning Based Java?

- A modeling language for learning and inference
- Supports
  - Programming using learned models

What is Learning Based Java?

- A modeling language for learning and inference
- Supports
  - Programming using learned models
  - High level specification of features and constraints between classifiers

What is Learning Based Java?

- A modeling language for learning and inference
- Supports
  - Programming using learned models
  - High level specification of features and constraints between classifiers
  - Inference with constraints

What is Learning Based Java?

- A modeling language for learning and inference
- Supports
  - Programming using learned models
  - High level specification of features and constraints between classifiers
  - Inference with constraints
- Learning
  - Classifiers are functions defined in terms of data

What is Learning Based Java?

- A modeling language for learning and inference
- Supports
  - Programming using learned models
  - High level specification of features and constraints between classifiers
  - Inference with constraints
- Learning
  - Classifiers are functions defined in terms of data
  - Learning can happen at compile time

What does LBJava do for you?

What does LBJava do for you?

- Abstracts away the feature representation, learning and inference

What does LBJava do for you?

- Abstracts away the feature representation, learning and inference

- Allows you to write learning based programs

What does LBJava do for you?

- Abstracts away the feature representation, learning and inference
- Allows you to write learning based programs
- Application developers can reason about the application at hand

Demo1: The Badges game

+ Naoki Abe	- Eric Baum

- Conference attendees to the 1994 Machine Learning conference were given name badges labeled with + or −.

- What function was used to assign these labels?

Why use learning?

- We typically use machine learning when the function f(**x**) we want the system to apply is too complex to program by hand.

Demo1: What's X for the Badges game?

- Possible features:
  - Gender/age/country of the person?
  - Length of their first or last name?
  - Does the name contain letter 'x'?
  - How many vowels does their name contain?
  - Is the n-th letter a vowel?

- Model this in LBJava, using the following features:
  - use the type of the characters in the first 5 positions of name
  - use the type of the characters in first 5 positions of the family name.

Demo1: What's X for the Badges game?

- Possible features:
  - Gender/age/country of the person?
  - Length of their first or last name?
  - Does the name contain letter 'x'?
  - How many vowels does their name contain?
  - Is the n-th letter a vowel?

- Model this in LBJava, using the following features:

me

mily

For example:
first-character-of-first-name-is-a
first-character-of-first-name-is-b …
second-character-of-first-name-is-a, …

Running on linux machine

Step 1: Compile Java code (Readers etc.)

- Need Java version 7 or higher

```
$ javac -cp “lib/*” -d bin *.java
```

Step 2: Compile (and train) the LBJava code

```
$ java -cp “lib/*:bin”
edu.illinois.cs.cogcomp.lbjava.Main -d bin
classifier.lbj
```

Demo2: Spam/noSpam

- The spam classifier
  1. The features
  2. The classifier
  3. Compiling to train the classifier

Demo2: Spam/noSpam

Demo2: Spam/noSpam

How a spam looks like? Features!

Demo2: Spam/noSpam

How a spam looks like? Features!

- Let us simply use features based on occurring words or maybe word frequencies.

Demo2: Spam/noSpam

How a spam looks like? Features!

- Let us simply use features based on occurring words or maybe word frequencies.
- Write our features and learners using Lbjava.

Demo3: Prediction of Drug Response for Cancer Patients

Input

Patient name
age, race, …

Patient name
gene1_Experimental result
gene2_Experimental result
…
genen_Experimental result

Output

If Patient X will response to Drug Y

Drug response is measured and reported as a real value but we can use a threshold and convert it to a binary decision of positive and negative response here.

Exercise

- Tweeter sentiment classification
  - http://l2r.cs.uiuc.edu/~danr/Teaching/CS446-15/readme-twitter.txt
- Train a classifier on annotated examples
- Predict sentiment of tweets in real time!
  - Filter by location, search terms, language, etc.

Links

- LBJava Software:

http://cogcomp.cs.illinois.edu/page/software_view/LBJava

- LBJava Manual:

http://cogcomp.cs.illinois.edu/software/manuals/LBJ2Manual.pdf

- Tutorial 2013 code and examples, step by step :

http://cogcomp.cs.illinois.edu/page/tutorial.201310

See you next time!

See you next time!

Parameter tuning

See you next time!

Parameter tuning

Designing more complex models

See you next time!

Parameter tuning

Designing more complex models

Pipelines

See you next time!

Parameter tuning

Designing more complex models

Pipelines

Inference and Constraints