Annotations, Evaluation, and Generalization

CS 485, Fall 2023
Applications of Natural Language Processing
https://people.cs.umass.edu/~brenocon/cs485_f23/

Brendan O'Connor
College of Information and Computer Sciences
University of Massachusetts Amherst

- Annotations
  - Chance-adjusted calculation
    - Practical ethics example in ChatGPT
- Evaluation
  - Held-out data and overfitting
  - Classification metrics
  - Statistical testing (J&M 4.9) - hold off until later in course

Cohen's Kappa for IAA

- If some classes predominate, raw agreement rate may be misleading
- Idea: normalize accuracy (agreement) rate such that answering randomly = 0.
  - From psychology / psychometrics / content analysis
- **Chance-adjusted agreement:**

$$\kappa = \frac{p_o - p_e}{1 - p_e}$$

p_o: **o**bserved agreement rate

p_e: **e**xpected (by chance) rate

Other chanced-adjusted metrics: Fleiss, Krippendorff... see reading

When is annotating ethical?

- Annotate stressful / traumatic content
- Confidentiality / privacy
- Transparency
- Pay !

Human labeling is key to ChatGPT

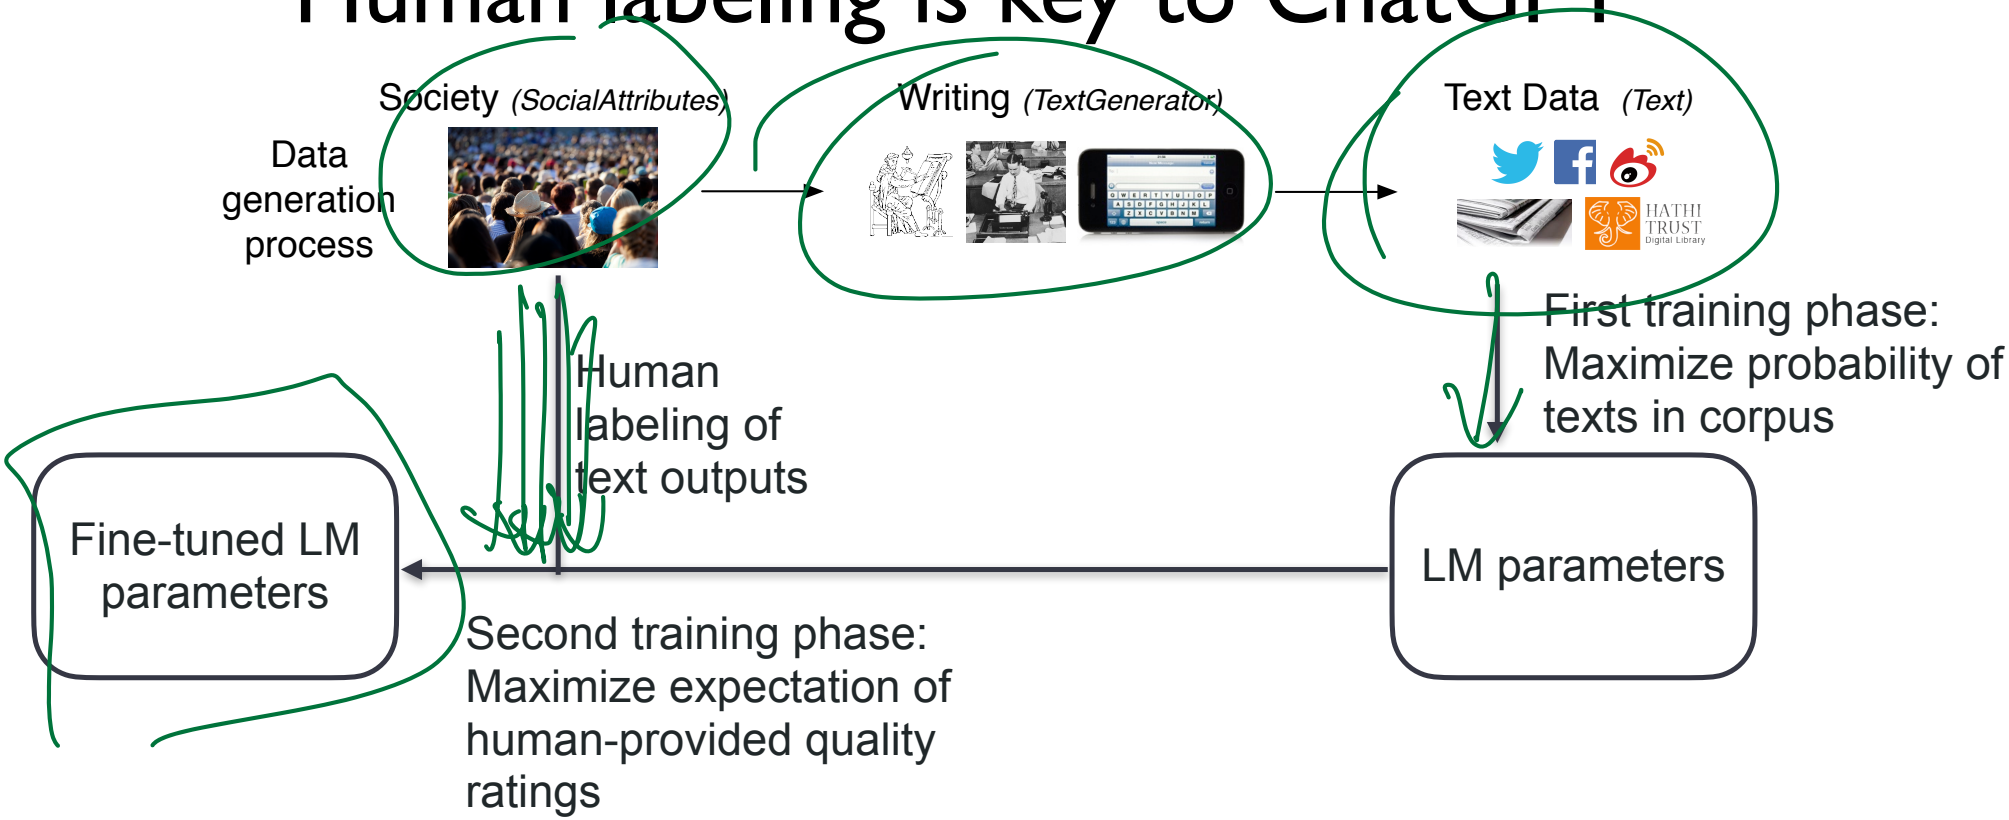

[Ouyang et al., 2022, Taori et al. 2023]

Table 3: Labeler-collected metadata on the API distribution.

Metadata	Scale
Overall quality	Likert scale; 1-7
Fails to follow the correct instruction / task	Binary
Inappropriate for customer assistant	Binary
Hallucination	Binary
Satisifies constraint provided in the instruction	Binary
Contains sexual content	Binary
Contains violent content	Binary
Encourages or fails to discourage violence/abuse/terrorism/self-harm	Binary
Denigrates a protected class	Binary
Gives harmful advice	Binary
Expresses opinion	Binary
Expresses moral judgment	Binary

[Ouyang et al., 2022]

'That Was Torture;' OpenAI Reportedly Relied on Low-Paid Kenyan Laborers to Sift Through Horrific Content to Make ChatGPT Palatable

The laborers reportedly looked through graphic accounts of child sexual abuse, murder, torture, suicide, and, incest.

By **Mack DeGeurin** Published January 18, 2023 | Comments (6) | Alerts

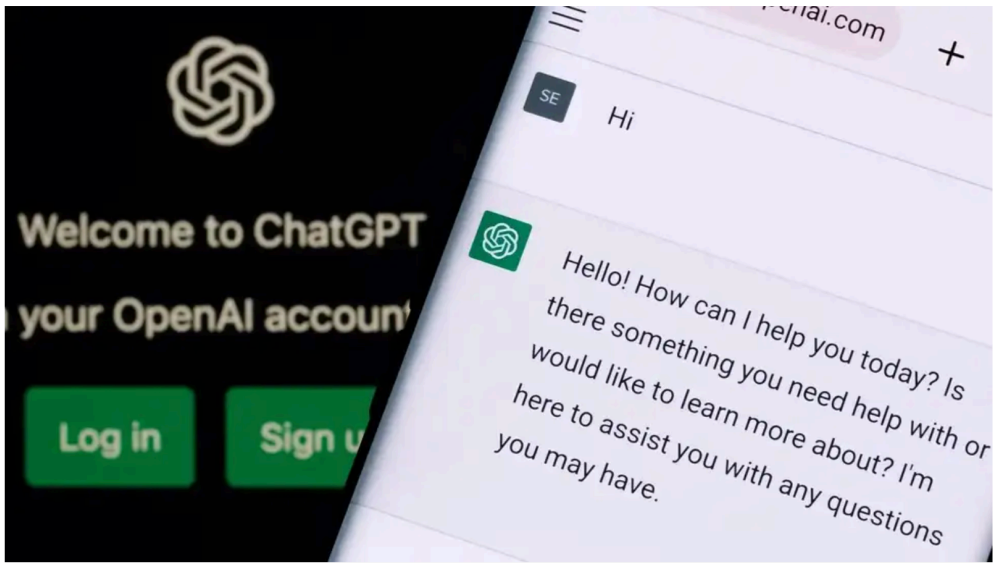

Image: Ascannio (Shutterstock)

Held-out data for evaluation

- How well will my classifier work in the future?
  - Can we look at classifier accuracy on training data?

x	y	p(y=1|x)
I love dog	1	.8
good dog	1	.7
cat cat	0	.3

"dog" occurs in $y=1$

never occurs in $y=0$

$$P(y=1|x;\beta) = \frac{1}{1+e^{-\beta^T x}} = g(\beta^T x)$$

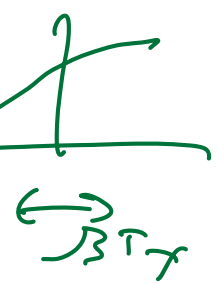

β_{dog} ↑ $\Rightarrow$ $\beta^T x = \sum_j \beta_j x_j$

$= \beta_{dog} x_{dog} + \beta_{101} x_{101} + \dots\dots$

Held-out data for evaluation

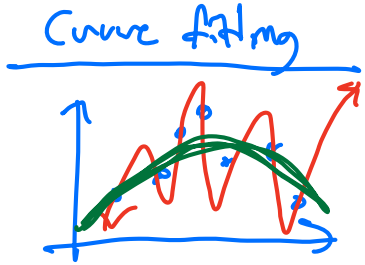

- Need to diagnose how much your model is **overfitting** the training set
- Data splits are key. Some ways to split:
  - Training set -vs- test set
  - Training set -vs- "validation"/"development" set -vs- test set
  - Cross-validation (within training set) -vs- test set

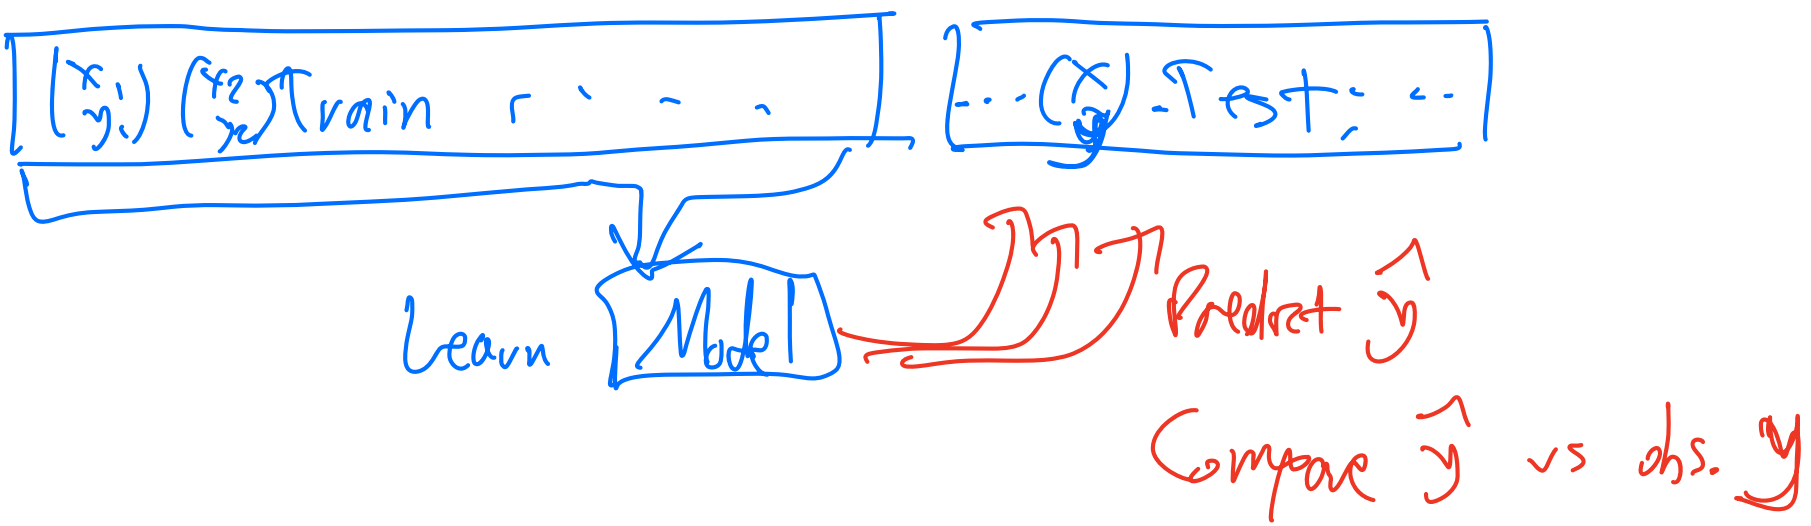

Cross-validation

- Cross-validation (within training set) -vs- test set

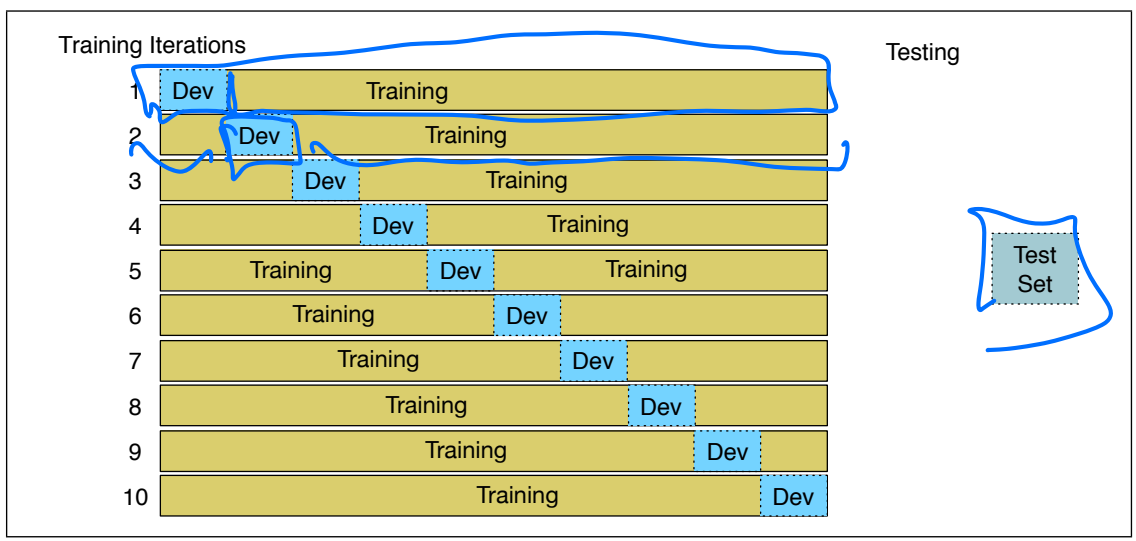

Figure 4.7 10-fold cross-validation

Regularization in Naive Bayes

$$\frac{p(y=1|\vec{w})}{p(y=0|\vec{w})} = \frac{p(y=1)\prod_i p(w_i|y=1)}{p(y=0)\prod_i p(w_i|y=0)}$$

$$= \frac{p(y=1)}{p(y=0)} \prod_i \frac{p(w_i|y=1)}{p(w_i|y=0)}$$

$LR(w)$

$$LR(dog) = \frac{[c(dog, y=1) + \alpha] / [c(y=1) + \alpha|V|]}{[c(dog, y=0) + \alpha] / [c(y=0) + \alpha|V|]}$$

$=0$

ignore

$\uparrow \alpha \Rightarrow$ higher regul. more biased model

$LR > 1$: favors $y=1$

$LR < 1$: favors $y=0$

- $\alpha = 0$ "
- $\alpha =$ tiny $\rightarrow LR \uparrow$
- $\alpha \to \infty \Rightarrow LR = 1$

Regularization in logistic regression

- If "dog" only occurs for class **k**, what weight will it get?
- Consider MLE training:

$$\arg\max_{\beta} \log P(y_1 \ldots y_n | x_1 \ldots x_n, \beta)$$

- Solution: ***regularized*** training for logistic regression

$$\arg\max_{\beta} \left[\log P(y_1 \ldots y_n | x_1 \ldots x_n) - \lambda \sum_{j=1}^{N_{feat}} (\beta_j)^2 \right]$$

Regularization tradeoffs

- No regularization <----------> Very strong regularization

Bad generalization

Great training acc!

Similar preds, on all examples

Similar test & tr. acc

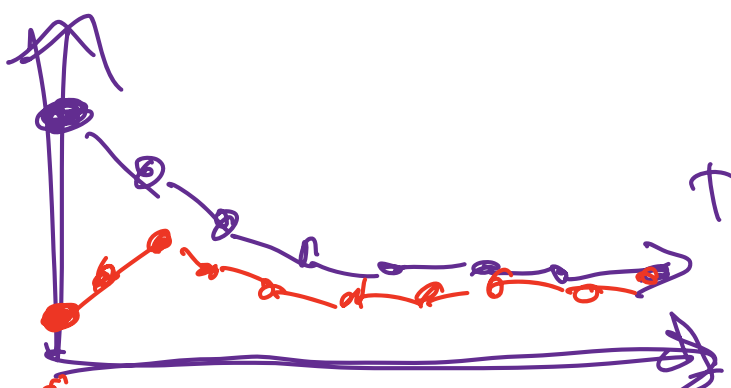

↑ tr acc ↑ Te acc

overfitting

or validation/dev data

Do I have enough labels?

- For training, hundreds to thousands of annotations may be needed for reasonable performance
  - Current work: how to usefully make NLP models with <10 or <100 training examples. "Few-shot learning"
- Exact amounts are difficult to know in advance. Can do a **learning curve** to estimate if more annotations will be useful.

Evaluation metrics

- Accuracy = $\frac{1}{n}\sum_{i=1}^{n} \mathbb{1}(\hat{y}^{(i)} = y^{(i)})$

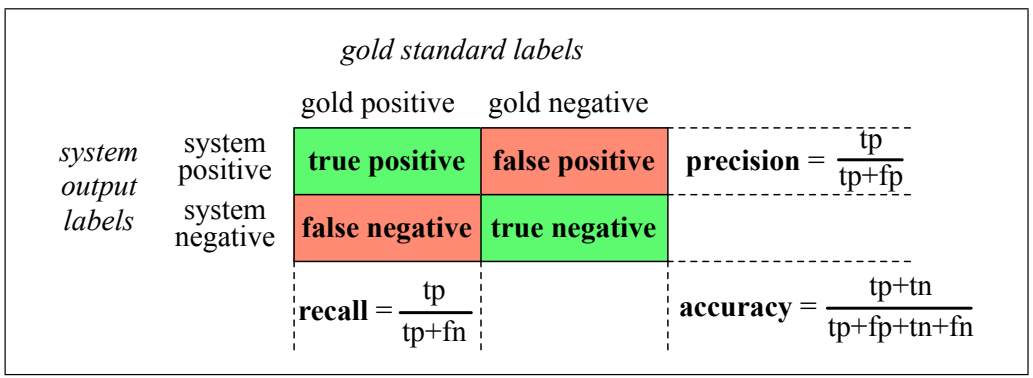

Figure 4.4 A confusion matrix for visualizing how well a binary classification system performs against gold standard labels.

- But do we care about false positives and negatives equally?
- What about rare classes?

Precision, recall, F1

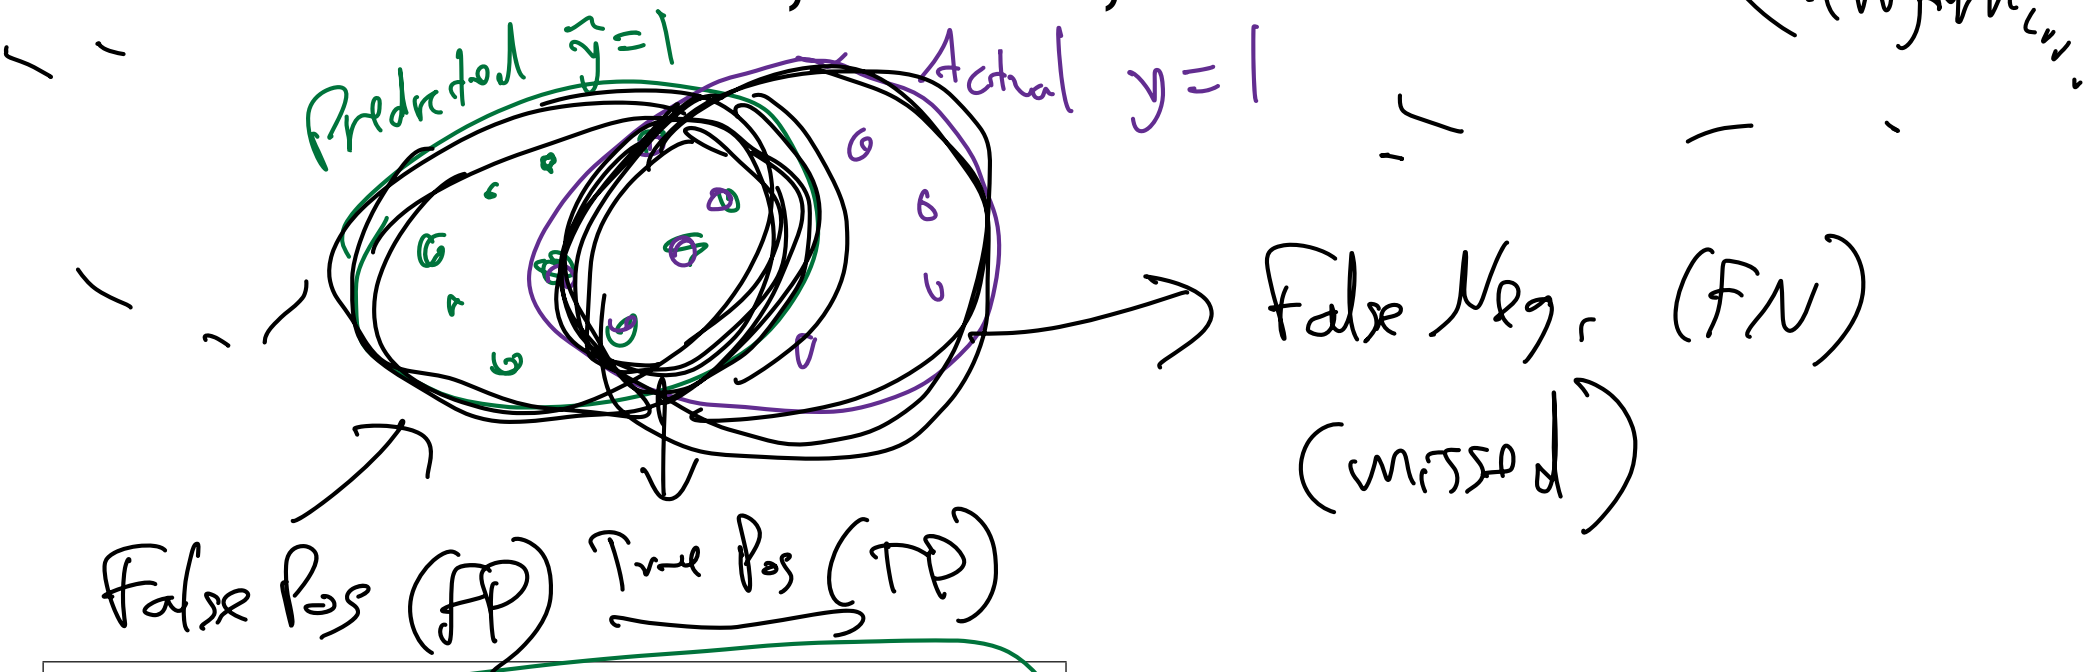

gold standard labels

system output labels		gold positive	gold negative	
	system positive	**true positive**	**false positive**	**precision** = $\frac{tp}{tp+fp}$
	system negative	**false negative**	**true negative**	
		recall = $\frac{tp}{tp+fn}$		**accuracy** = $\frac{tp+tn}{tp+fp+tn+fn}$

Figure 4.4 A confusion matrix for visualizing how well a binary classification system performs against gold standard labels.

- macro vs. micro F1