

Stacking With Auxiliary Features (SWAF)

Nazneen Fatema Rajani and Ray Mooney

nrajani@cs.utexas.edu and mooney@cs.utexas.edu

Dept. of Computer Science at the University of Texas at Austin

INTRODUCTION

- Stacking (Wolpert, 1992) is a well known **ensembling** algorithm
- However, it does not adequately discriminate between **base systems** and **input instances**
- Stacking With Auxiliary Features (SWAF) integrates information from multiple sources
- Auxiliary Features enable the stacker to leverage relevant information to improve prediction
- We use two types of auxiliary features :
 - Instance features – enable the stacker to discriminate across instances
 - Provenance features – enable the stacker to discriminate across base systems

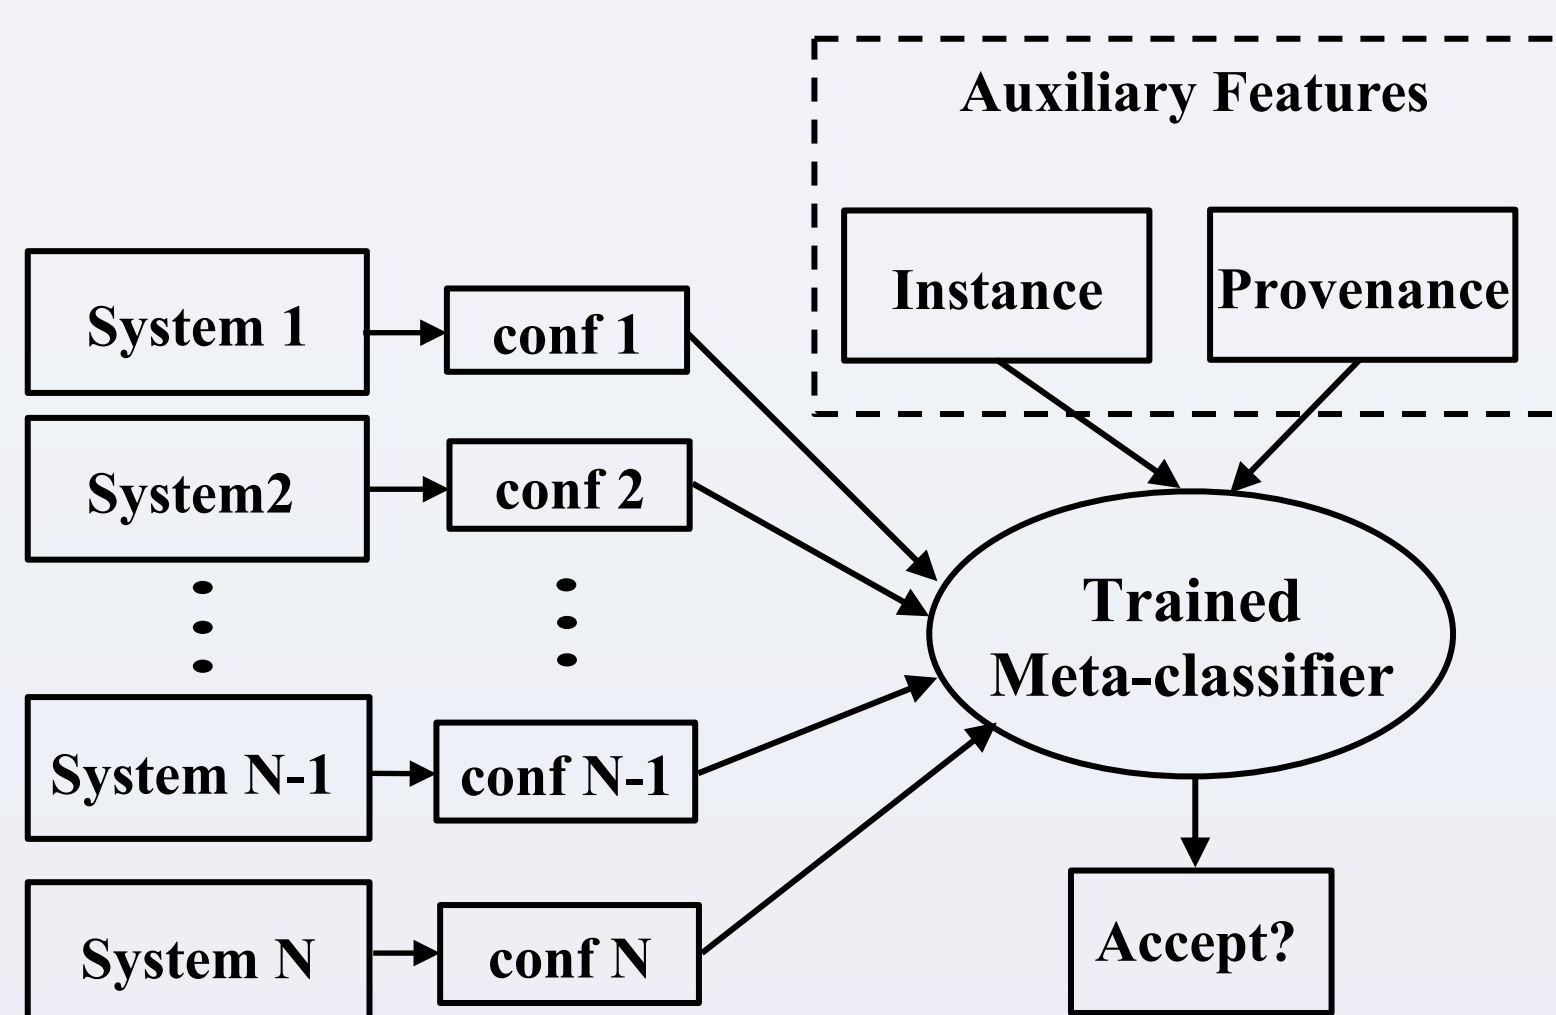


Figure 1: Our stacking approach to combining system outputs using confidence scores and two types of auxiliary features for improving prediction

TASK OVERVIEW

- We demonstrate SWAF on three very different machine learning problems
- Two of them are in NLP and third is a well known computer vision problem

Slot Filling (SF)

org: Microsoft	
1. city_of_headquarters:	Microsoft is a technology company, headquartered in Redmond, Washington that develops ...
2. website:	
3. subsidiaries:	city_of_headquarters: Redmond
4. employees:	provenance: [redacted]
5. shareholders:	confidence score: 1.0
⋮	

Figure 2: SF involves building a Knowledge Base (KB) from scratch using pre-defined slots. Systems provide confidence score and provenance

Entity Discovery and Linking (EDL)

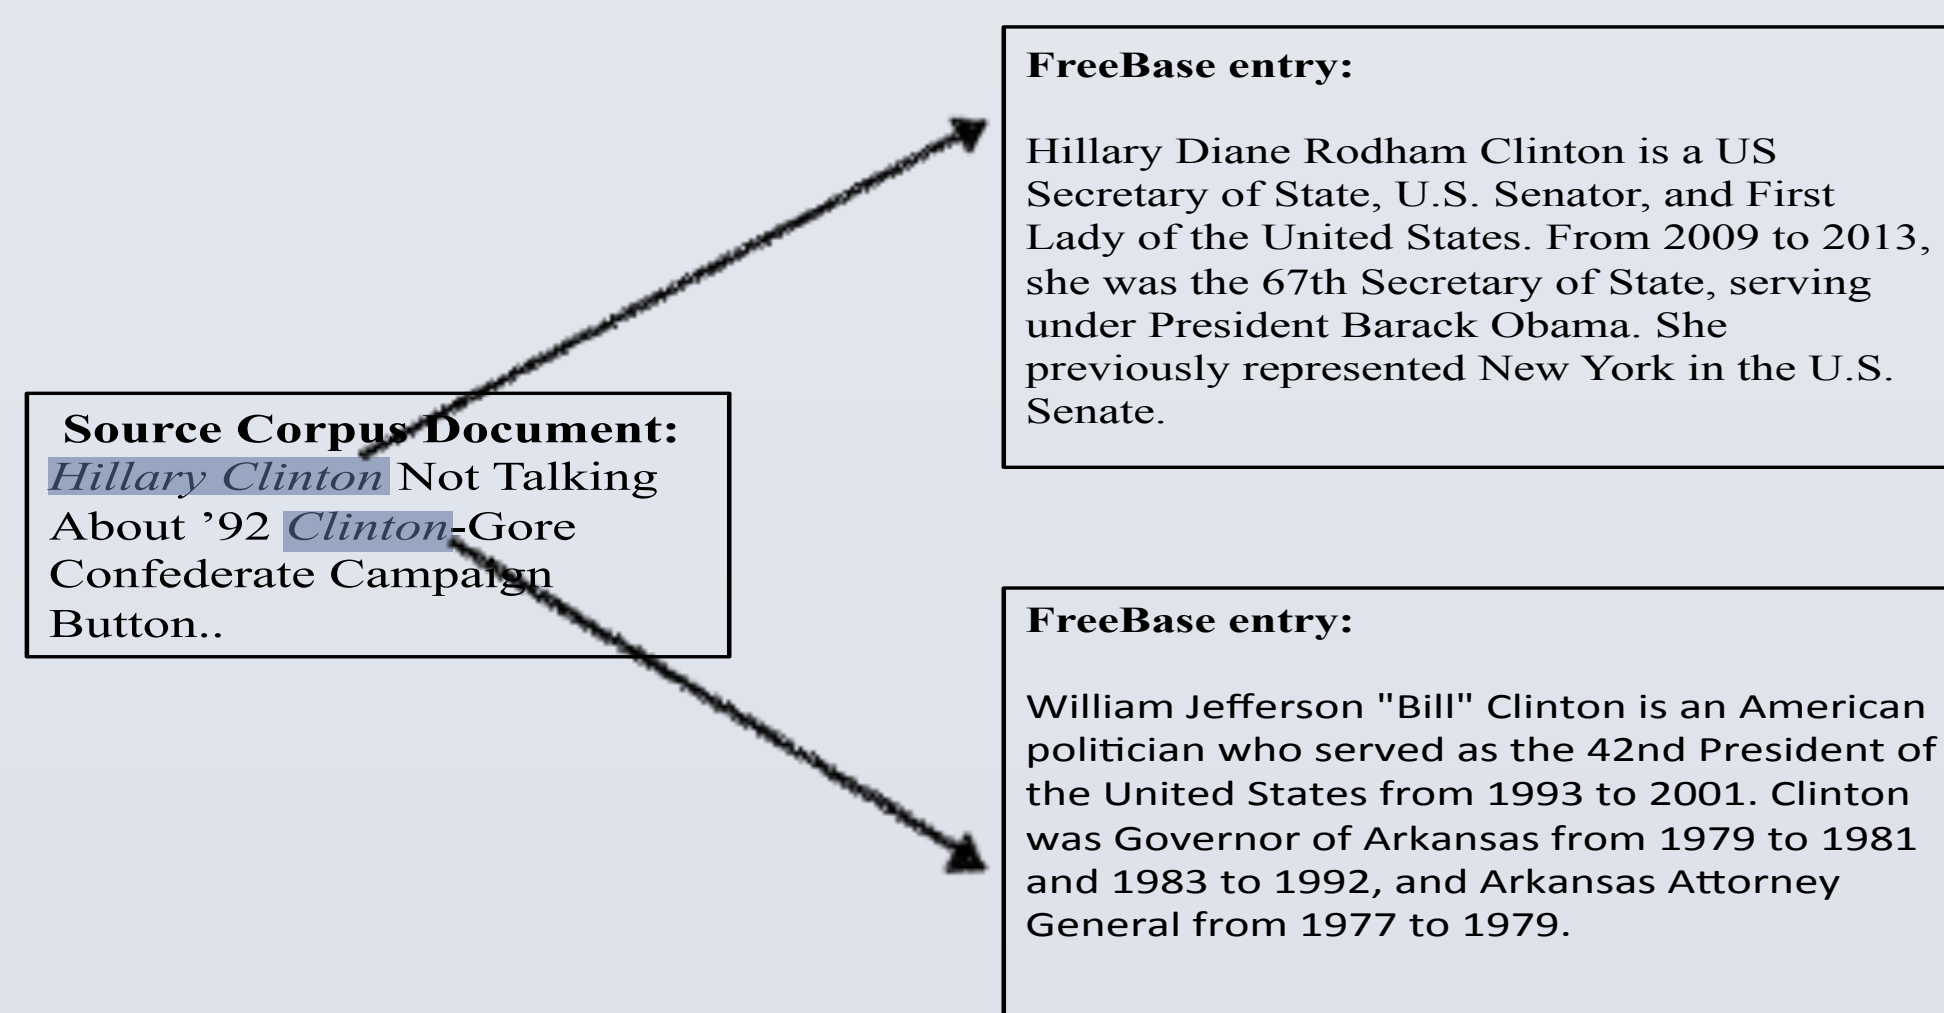


Figure 3: EDL involves detecting entity mentions in a corpus and linking them to an English KB (FreeBase). Systems provide confidence score and provenance

ImageNet Object Detection

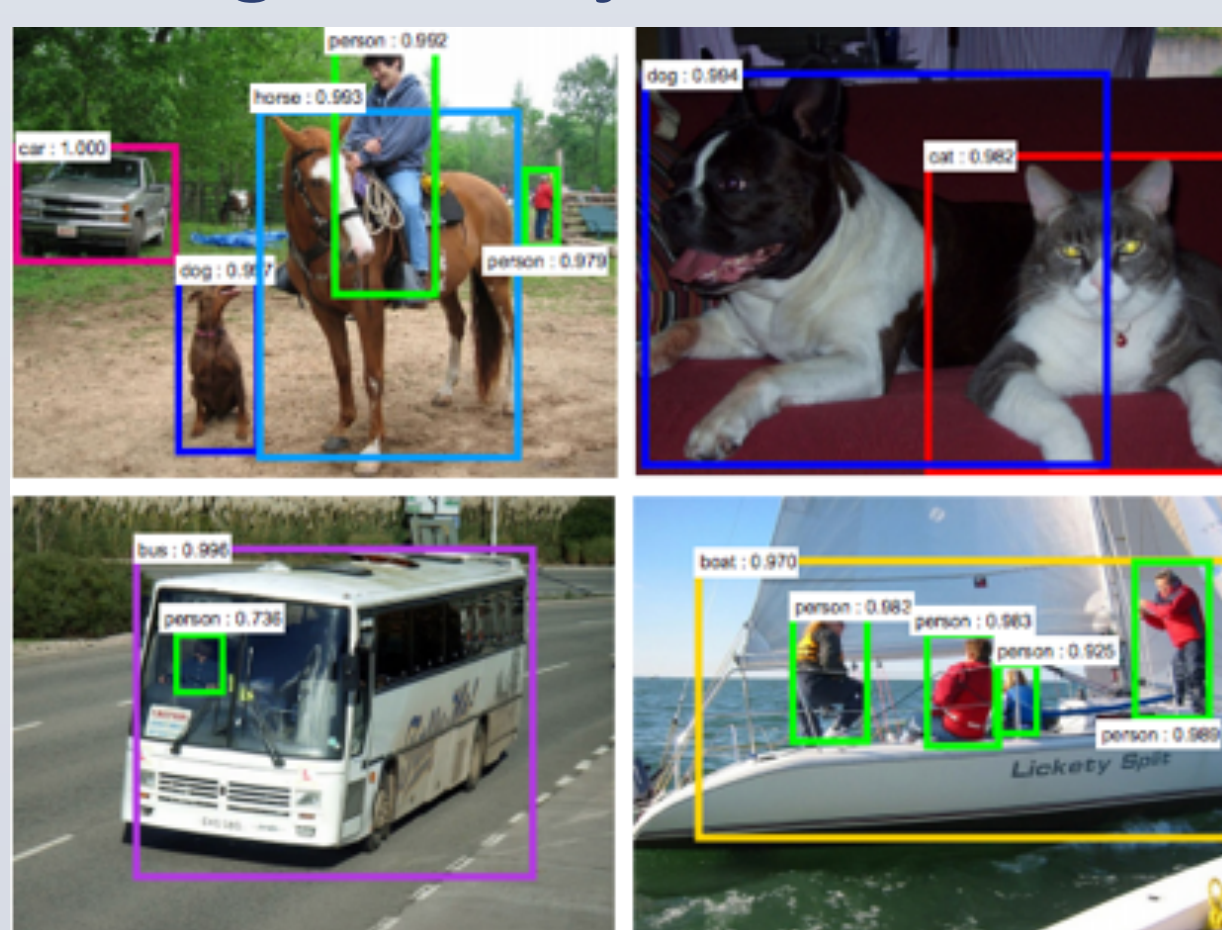


Figure 4: Detect all instances of object categories (total 200) in images and localize using bounding boxes

AUXILIARY FEATURES


Instance Features

- Enables the stacker to discriminate between **input instance types**
- The intuition is that some systems are **better** at certain inputs than other systems
- Information** about the input type would allow the classifier to make a better prediction
- Slot Filling – slot type (per: age, org: headquarters)
- Entity Detection and Linking – entity type (PER, ORG, GPE)
- Object Detection – object category (200) and VGGNet's *fc7* features

Provenance Features


- Enables the stacker to discriminate between **component systems**
- The intuition is that output is **reliable** if systems agree on the source or provenance
- Information** about provenance or source of the output would allow the classifier to make a better prediction
- Slot Filling :-
 - Document provenance – For a given query and slot, for each system, i , there is a feature DP_i :
 - N systems provide a fill for the slot.
 - Of these, n give same provenance *docid* as i .
 - $DP_i = n/N$ is the document provenance score.
 - Offset provenance - Degree of overlap between systems' provenance strings. Uses Jaccard similarity coefficient.

Offsets	System 1	System 2	System 3
Start Offset	1	4	5
End Offset	9	7	12



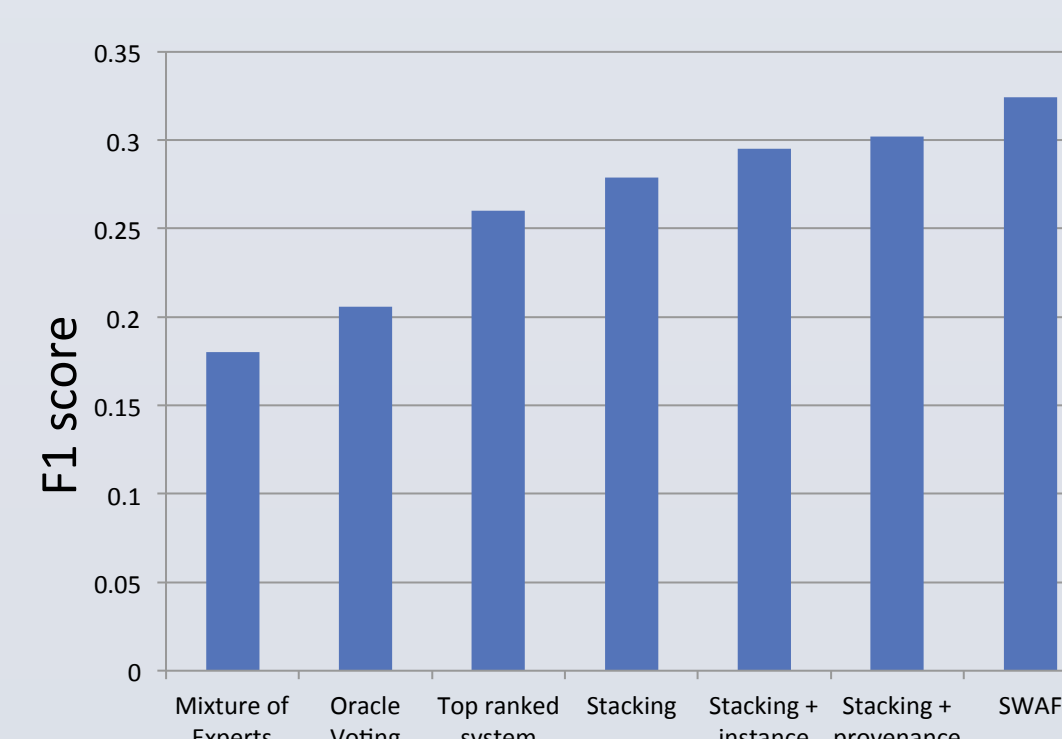
$$OP_1 = \frac{1}{2} \times \left(\frac{4}{9} + \frac{5}{12} \right)$$

- Entity Detection and Linking (EDL) :- Same as Slot Filling using the entity mention as provenance
- Object Detection :- Bounding box overlap measured using Jaccard similarity coefficient



RESULTS

Slot Filling



Entity Discovery and Linking



ImageNet Object Detection

