

Named Entity Recognition

Asif Ekbal

Computer Science and Engineering

IIT Patna, India-800 013

Email: asif@iitp.ac.in

asif.ekbal@gmail.com

IASNLP 2014

Outline

- Background
- Introduction to the various issues of NER
- NER in different languages
- NER in Indian languages
- Evolutionary Approaches to NER
 - Brief Discussions on Genetic Algorithm
 - Some Issues of Classifier Ensemble

Outline

- Single Objective Optimization
 - Weighted Vote based Classifier Ensemble

- Multiobjective Optimization
 - Brief introduction to MOO
 - Classifier Ensemble Selection

- Bio-Text Mining
 - Introduction
 - NE Extraction in Biomedicine

Background: Information Extraction

- To extract information that fits **pre-defined** database schemas or templates, specifying the output formats
- **IE Definition**
 - **Entity**: An object of interest such as a person or organization
 - **Attribute**: A property of an entity such as name, alias, descriptor or type
 - **Fact**: A relationship held between two or more entities such as Position of Person in Company
 - **Event**: An activity involving several entities such as terrorist act, airline crash, product information

The Problem

The diagram illustrates the process of extracting structured information from a text document. It features a central text block with several lines of text. Five orange callout boxes with black outlines point to specific parts of the text, each containing a label. The labels are: 'Date' (pointing to 'Friday, March 24, 2006'), 'Time: Start - End' (pointing to '9:30-11:00 a.m.'), 'Location' (pointing to '1014 DOW'), 'Speaker' (pointing to 'Dave Lewis'), and 'Person' (pointing to 'Gady Agam, Shlomo Argamon, Ophir Frieder, Dave Grossman').

DATE: Friday, March 24, 2006

TIME: 9:30-11:00 a.m.

LOCATION: 1014 DOW

SPEAKER: Dave Lewis

TITLE: Bayesian Logistic Regression Classification and Mining (Plus A Big New Test Collection)

ABSTRACT

Bayesian logistic regression allows incorporating task knowledge through model structure and priors on parameters. I will discuss content-based text categorization and authorship attribution using 1) priors that control sparsity and sign of parameters, 2) priors that incorporate domain knowledge from reference books and other texts, and 3) the use of polytomous (1-of-k) dependent variables. All experiments were performed with our open-source programs, BBR and BMR, which can fit models with millions of parameters. (Joint work with David Madigan, Alex Genkin, Aynur Dayanik, Dmitriy Fradkin, and Vladimir Menkov at Rutgers and DIMACS.) I will also briefly discuss the IIT CDIP (Complex Document Information Processing) test collection, which I am developing under an ARDA subcontract to Illinois Institute of Technology. It is based on 1.5TB of scanned and OCR'd documents released in tobacco litigation, and will be a major resource for research in information retrieval, document analysis, social network analysis, and perhaps databases. (Joint work with Gady Agam, Shlomo Argamon, Ophir Frieder, Dave Grossman, and David Madigan.)

BIOGRAPHY

Dave Lewis is based in Chicago, IL, and consults on information retrieval, data mining, and natural language processing. He previously held research positions at AT&T Labs, Bell Labs, and the University of Chicago. He received his Ph.D. in Computer Science from the University of Massachusetts, Amherst, and did his undergraduate work down the road at Michigan State.

What is “Information Extraction”?

As a task: Filling slots in a database from sub-segments of text.

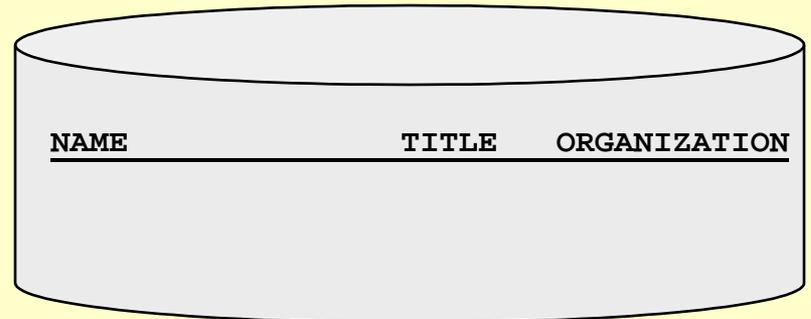
October 14, 2002, 4:00 a.m. PT

For years, Microsoft Corporation CEO Bill Gates railed against the economic philosophy of open-source software with Orwellian fervor, denouncing its communal licensing as a "cancer" that stifled technological innovation.

Today, Microsoft claims to "love" the open-source concept, by which software code is made public to encourage improvement and development by outside programmers. Gates himself says Microsoft will gladly disclose its crown jewels--the coveted code behind the Windows operating system--to select customers.

"We can be open source. We love the concept of shared source," said Bill Veghte, a Microsoft VP. "That's a super-important shift for us in terms of code access."

Richard Stallman, founder of the Free Software Foundation, countered saying...



What is “Information Extraction”

As a task:

Filling slots in a database from sub-segments of text

October 14, 2002, 4:00 a.m. PT

For years, Microsoft Corporation **CEO** Bill Gates railed against the economic philosophy of open-source software with Orwellian fervor, denouncing its communal licensing as a "cancer" that stifled technological innovation.

Today, Microsoft claims to "love" the open-source concept, by which software code is made public to encourage improvement and development by outside programmers. Gates himself says Microsoft will gladly disclose its crown jewels--the coveted code behind the Windows operating system--to select customers.

"We can be open source. We love the concept of shared source," said Bill Veghte, a Microsoft **VP**. "That's a super-important shift for us in terms of code access."

Richard Stallman, **founder** of the Free Software Foundation, countered saying...



| <u>NAME</u> | <u>TITLE</u> | <u>ORGANIZATION</u> |
|------------------|--------------|---------------------|
| Bill Gates | CEO | Microsoft |
| Bill Veghte | VP | Microsoft |
| Richard Stallman | founder | Free Soft.. |

What is “Information Extraction”

Information Extraction =
segmentation + classification + association + clustering

October 14, 2002, 4:00 a.m. PT

For years, Microsoft Corporation CEO Bill Gates railed against the economic philosophy of open-source software with Orwellian fervor, denouncing its communal licensing as a "cancer" that stifled technological innovation.

Today, Microsoft claims to "love" the open-source concept, by which software code is made public to encourage improvement and development by outside programmers. Gates himself says Microsoft will gladly disclose its crown jewels--the coveted code behind the Windows operating system--to select customers.

"We can be open source. We love the concept of shared source," said Bill Veghte, a Microsoft VP. "That's a super-important shift for us in terms of code access."

Richard Stallman, founder of the Free Software Foundation, countered saying...

Microsoft Corporation

CEO

Bill Gates

Microsoft

Gates

Microsoft

Bill Veghte

Microsoft

VP

Richard Stallman

founder

Free Software Foundation

**aka “named entity
extraction”**

What is “Information Extraction”

Information Extraction =
segmentation + classification + association + clustering

October 14, 2002, 4:00 a.m. PT

For years, [Microsoft Corporation](#) [CEO](#) [Bill Gates](#) railed against the economic philosophy of open-source software with Orwellian fervor, denouncing its communal licensing as a "cancer" that stifled technological innovation.

Today, Microsoft claims to "love" the open-source concept, by which software code is made public to encourage improvement and development by outside programmers. Gates himself says Microsoft will gladly disclose its crown jewels--the coveted code behind the Windows operating system--to select customers.

"We can be open source. We love the concept of shared source," said [Bill Veghte](#), a [Microsoft](#) [VP](#). "That's a super-important shift for us in terms of code access."

[Richard Stallman](#), [founder](#) of the [Free Software Foundation](#), countered saying...

Microsoft Corporation

CEO

Bill Gates

Microsoft

Gates

Microsoft

Bill Veghte

Microsoft

VP

Richard Stallman

founder

Free Software Foundation

What is “Information Extraction”

**Information Extraction =
segmentation + classification + association + clustering**

October 14, 2002, 4:00 a.m. PT

For years, Microsoft Corporation **CEO** Bill Gates railed against the economic philosophy of open-source software with Orwellian fervor, denouncing its communal licensing as a "cancer" that stifled technological innovation.

Today, Microsoft claims to "love" the open-source concept, by which software code is made public to encourage improvement and development by outside programmers. Gates himself says Microsoft will gladly disclose its crown jewels--the coveted code behind the Windows operating system--to select customers.

"We can be open source. We love the concept of shared source," said Bill Veghte, a Microsoft **VP**. "That's a super-important shift for us in terms of code access."

Richard Stallman, **founder** of the Free Software Foundation, countered saying...

Microsoft Corporation
CEO
Bill Gates

Microsoft
Gates

Microsoft
Bill Veghte
Microsoft
VP

Richard Stallman
founder
Free Software Foundation

What is "Information Extraction"

**Information Extraction =
segmentation + classification + association + clustering**

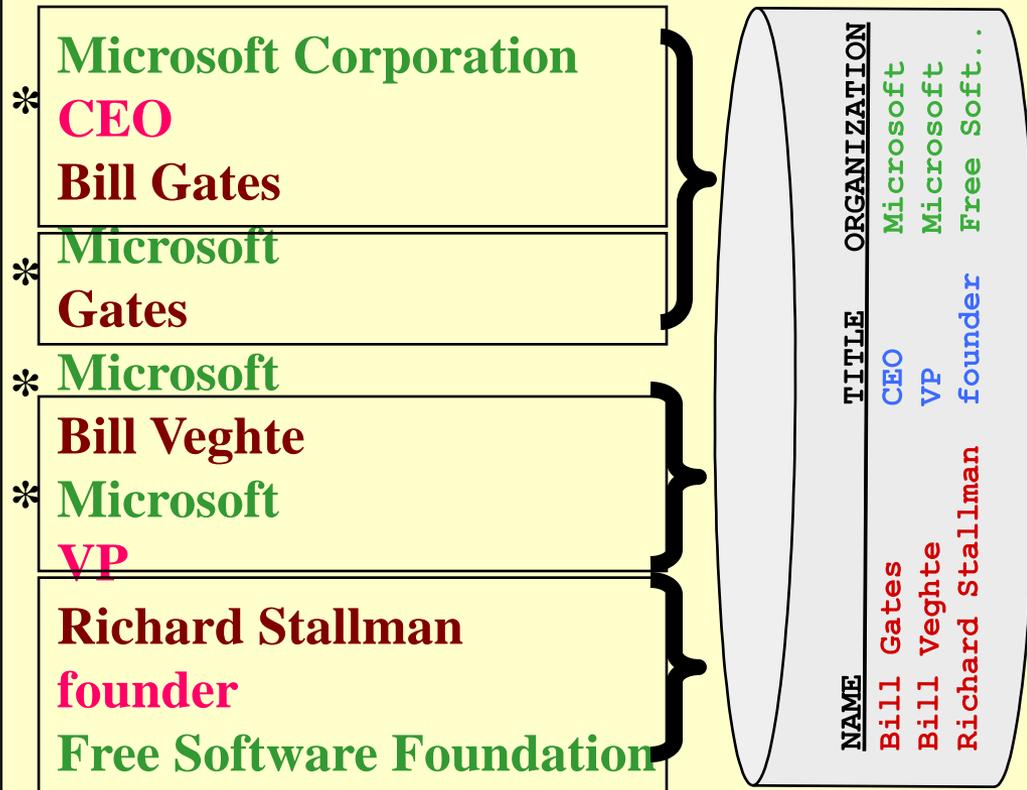
October 14, 2002, 4:00 a.m. PT

For years, [Microsoft Corporation](#) [CEO](#) [Bill Gates](#) railed against the economic philosophy of open-source software with Orwellian fervor, denouncing its communal licensing as a "cancer" that stifled technological innovation.

Today, Microsoft claims to "love" the open-source concept, by which software code is made public to encourage improvement and development by outside programmers. Gates himself says Microsoft will gladly disclose its crown jewels--the coveted code behind the Windows operating system--to select customers.

"We can be open source. We love the concept of shared source," said [Bill Veghte](#), a [Microsoft](#) [VP](#). "That's a super-important shift for us in terms of code access."

[Richard Stallman](#), [founder](#) of the [Free Software Foundation](#), countered saying...



What is Named Entity Recognition?

□ **NER – Named Entity Recognition** involves identification of proper names in texts, and classification into a set of pre-defined categories of interest as:

- **Person names** (names of people)
- **Organization names** (companies, government organizations, committees, etc.)
- **Location names** (cities, countries etc)
- **Miscellaneous names** (Date, time, number, percentage, monetary expressions, number expressions and measurement expressions)

Named Entity Recognition

Markables (as defined in MUC6 and MUC7)

Names of **organization**, **person**, **location**

Mentions of **date** and **time**, **money** and **percentage**

Example:

“Ms. **Washington**'s candidacy is being championed by several powerful lawmakers including her boss, Chairman **John Dingell** (D., **Mich.**) of the **House Energy and Commerce Committee.**”

Task Definition

- **Other common types**: measures (percent, money, weight etc), email addresses, web addresses, street addresses, etc.
- **Some domain-specific entities**: names of drugs, medical conditions, names of ships, bibliographic references etc.
- MUC-7 entity definition guidelines (Chinchor'97)

http://www.itl.nist.gov/iaui/894.02/related_projects/muc/proceedings/ne_task.html

Basic Problems in NER

- Generative in nature
- Variation of NERs – e.g. Prof Manning, Chris Manning, Dr Chris Manning
- Ambiguity of NE types:
 - **Washington** (location vs. person)
 - **May** (person vs. month)
 - **Ford** (person vs organization)
 - **1945** (date vs. time)
- Ambiguity with common words, e.g. “*Kabita*”
 - Name of person vs. poem

More complex problems in NER

- Issues of style, structure, domain, genre etc.
- Punctuation, spelling, spacing, formatting, ... all have an impact:

Dept. of Computing and Maths

Manchester Metropolitan University

Manchester

United Kingdom

Applications

- **Intelligent document access**
 - Browse document collections by the entities that occur in them
 - Application domains:
 - News
 - Scientific articles, e.g, MEDLINE abstracts
- **Information retrieval and extraction**
 - Augmenting a query given to a retrieval system with NE information, more refined information extraction is possible
 - For example, if a person wants to search for document containing '*kabiTA*' as a proper noun, adding the NE information will eliminate irrelevant documents with only '*kabiTA*' as a common noun

Applications

- **Machine translation**

- NER plays an important role in translating documents from one language to other
- Often the NEs are transliterated rather than translated
- For example, ‘*yAdabpur bishvabidyAlaYa*’ → ‘Jadavpur University’

- **Automatic summarization**

- NEs given more priorities in deciding the summary of a text
- Paragraphs containing more NEs are most likely to be included into the summary

Applications

- **Question-Answering Systems**
 - NERs are important to retrieve the answers of particular questions
- **Speech Related Tasks**
 - In Text to Speech (TTS), NER is important for identifying the number format, telephone number and date format
 - In speech rhythm- necessary to provide a short break after the name of person
 - Solving Out Of Vocabulary words is important in speech recognition

Corpora, Annotation

Some NE Annotated Corpora

- MUC-6 and MUC-7 corpora - English
- CONLL shared task corpora
 - <http://cnts.uia.ac.be/conll2003/ner/>: NEs in English and German
 - <http://cnts.uia.ac.be/conll2002/ner/>: NEs in Spanish and Dutch
- ACE – English - <http://www ldc.upenn.edu/Projects/ACE/>
- TIDES surprise language exercise (NEs in Hindi and Cebuano) - 2003
- NERSSEAL shared task- NEs in Bengali, Hindi, Telugu, Oriya and Urdu (<http://ltrc.iiit.ac.in/ner-ssea-08/index.cgi?topic=5>)-2008

Corpora, Annotation

- Biomedical and Biochemical corpora
 - BioNLP-04 shared task
 - BioCreative shared tasks
 - AiMed

Performance Evaluation

- **Evaluation metric** – mathematically defines how to measure the system's performance against a human-annotated, gold standard
- **Scoring program**—implements the metric and provides performance measures
 - For each document and over the entire corpus
 - For each type of NE

The Evaluation Metric

Precision = correct answers/answers produced

Recall = correct answers/total possible correct answers

Trade-off between precision and recall

F-Measure = $(\beta^2 + 1)PR / \beta^2R + P$

β reflects the weighting between precision and recall,
typically $\beta=1$

The Evaluation Metric (2)

Precision =

$$\frac{\text{Correct} + \frac{1}{2} \text{ Partially correct}}{\text{Correct} + \text{Incorrect} + \text{Partial}}$$

Recall =

$$\frac{\text{Correct} + \frac{1}{2} \text{ Partially correct}}{\text{Correct} + \text{Missing} + \text{Partial}}$$

NE boundaries are often misplaced, so
some partially correct results

Named Entity Recognition

- Handcrafted systems
 - Knowledge (rule) based
 - Patterns
 - Gazetteers
- Automatic systems
 - Statistical
 - Machine learning-*Supervised, Semi-supervised, Unsupervised*
- Hybrid systems

Pre-processing for NER

- Format detection
- Word segmentation (for languages like Chinese)
- Tokenisation
- Sentence splitting
- Part-of-Speech (PoS) tagging

Comparisons between two Approaches

Knowledge Engineering

- rule based
- developed by experienced language engineers
- makes use of human intuition
- requires only small amount of training data
- development could be very time consuming
- some changes may be hard to accommodate

Learning Systems

- use statistics or other machine learning
- developers do not need LE expertise
- requires large amounts of annotated training data
- annotators are cheap (but you get what you pay for!)
- easily trainable and adaptable to new domains and languages

List lookup approach-baseline

- System that recognises only entities stored in its lists (gazetteers)
- **Advantages** - Simple, fast, language independent, easy to retarget (just create lists)
- **Disadvantages** - collection and maintenance of lists, cannot deal with name variants, cannot resolve ambiguity

Shallow (internal structure)

Parsing

Approach

- Internal evidence—names often have internal structure. These components can be either stored or guessed,

e.g. location:

- Cap. Word + {City, Forest, Centre, River}

e.g. Sundarban Forest

- Cap. Word + {Street, Boulevard, Avenue, Crescent, Road}

e.g. MG Road

e.g. Person

- Word + {Kumar, Chandra} + Word

E.g., Naresh Kumar Singh

Problems with the shallow parsing approach

- Ambiguously capitalized words (first word in sentence)
[All American Bank] vs. All [State Police]
- Semantic ambiguity
"John F. Kennedy" = airport (location) vs. person
"Philip Morris" = organization vs. person
- Structural ambiguity
[Cable and Wireless] vs. [Microsoft] and [Dell]
[Center for Computational Linguistics] vs. message from
[City Hospital] for [John Smith]

Shallow Parsing Approach with Context

- Use of context-based patterns is helpful in ambiguous cases
- "David Walton" and "Goldman Sachs" are indistinguishable
 - But with the phrase "David Walton of Goldman Sachs" and the Person entity "David Walton" recognised, we can use the pattern "[Person] of [Organization]" to identify "Goldman Sachs" correctly

Examples of context patterns

- [PERSON] *earns* [MONEY]
- [PERSON] *joined* [ORGANIZATION]
- [PERSON] *left* [ORGANIZATION]
- [PERSON] *joined* [ORGANIZATION] as [JOBTITLE]
- [ORGANIZATION]'s [JOBTITLE] [PERSON]
- [ORGANIZATION] [JOBTITLE] [PERSON]
- *the* [ORGANIZATION] [JOBTITLE]
- *part of the* [ORGANIZATION]
- [ORGANIZATION] headquarters in [LOCATION]
- *price of* [ORGANIZATION]
- *sale of* [ORGANIZATION]
- *investors in* [ORGANIZATION]
- [ORGANIZATION] *is worth* [MONEY]
- [JOBTITLE] [PERSON]
- [PERSON], [JOBTITLE]

Named Entity Recognition

- Handcrafted systems
 - LTG (Mikheev et al., 1997)
 - F-measure of 93.39 in MUC-7 (the best)
 - Ltquery, XML internal representation
 - Tokenizer, POS-tagger, SGML transducer
 - Nominator (1997)
 - IBM
 - Heavy heuristics
 - Cross-document co-reference resolution
 - Used later in IBM Intelligent Miner

Named Entity Recognition

- Handcrafted systems
 - LaSIE (Large Scale Information Extraction)
 - MUC-6 (LaSIE II in MUC-7)
 - Univ. of Sheffield's GATE architecture (General Architecture for Text Engineering)
 - FACILE (1998)- Fast and Accurate Categorisation of Information by Language Engineering
 - NEA language (Named Entity Analysis)
 - Context-sensitive rules
 - NetOwl (MUC-7)
 - Commercial product
 - C++ engine, extraction rules

Gazetteer lists for rule-based NER

- Needed to store the indicator strings for the internal structure and context rules
- Internal location indicators – e.g., {*river, mountain, forest*} for natural locations; {*street, road, crescent, place, square, ...*} for address locations
- Internal organisation indicators—e.g., company designators {*GmbH, Ltd, Inc, ...*}
- Produces Lookup results of the given kind

Named Entities in GATE

Gate 2.1-alpha1 build 875

File Options Tools Help

Messages GATE corpus_0003D ANNIE_0001E example document

Text Annotations Annotation Sets Coreference Print

the business market, which is where local companies make a lot of their money."

The deal also will give three major cable television companies, which are majority owners of **Teleport**, a collective 10 percent stake in **AT&T**.

By acquiring **Teleport**, **AT&T** can offer business customers local and long-distance telephone service, and data and Internet access, under its own brand name. Using **Teleport**'s local facilities, the company also would be able to reduce the fees it pays to local phone companies for access to local telephone customers.

"It's going to permit us to be much more cost-effective as we go for that local business," Armstrong said at a news briefing. "This has competition and growth written all over it." **AT&T** is paying for **Teleport** with its stock. **Teleport** shareholders will receive 0.943 **AT&T** shares for each of their **Teleport** shares, putting the deal at \$ 59 a share based on **AT&T**'s closing price yesterday of \$ 62.62 1/2 a share, up \$ 2.62 1/2. **Teleport** closed down \$ 3.62 1/2 at \$ 54.12 1/2 a share. The companies expect the deal, which must be approved by regulators and shareholders, to close by fall. **Teleport**, based in Staten Island, N.Y., leads a new breed of local phone competitors that are invading urban markets to grab the most lucrative business customers from the regional Bell companies, **GTE Corp**.

| Type | Set | Start ▲ | End | |
|--------------|---------|---------|-----|--|
| Organization | Default | 37 | 52 | {rule2=OrgFinal, orgType=company, rul |
| Organization | Default | 115 | 130 | {rule2=OrgFinal, orgType=null, rule1=C |
| Organization | Default | 138 | 146 | {NMRule=Unknown, kind=PN, rule=Unk |
| Organization | Default | 148 | 161 | {rule2=OrgFinal, orgType=null, rule1=C |

Annotations Editor Features Editor

ANNIE_0001E run in 4.247 seconds

Default annotation

- Date
- FirstPerso
- JobTitle
- Location
- Lookup
- Money
- Organizati
- Percent
- Person
- Sentence
- SpaceTok
- Split
- Title
- Token
- Unknown

Original markups

- DOC
- DOCNO
- DOCTYPE
- HEADER

Using co-reference to classify ambiguous NEs

- Orthographic co-reference module that matches proper names in a document
- Improves NE results by assigning entity type to previously unclassified names, based on relations with classified NEs
- May not reclassify already classified entities
- Classification of unknown entities is very useful for surnames which match a full name, or abbreviations, e.g. [Bonfield] will match [Sir Peter Bonfield];
[International Business Machines Ltd.] will match [IBM]

Named Entity Coreference

The screenshot displays a software interface for text annotation. At the top, there are tabs for 'Messages', 'GATE corpus_0003D', 'ANNIE_0001E', and 'example document'. Below this is a navigation bar with buttons for 'Text', 'Annotations', 'Annotation Sets', 'Coreference', 'Print', and a magnifying glass icon. The main text area contains several paragraphs with named entities highlighted in colored boxes: 'Mike Armstrong' (blue), 'AT&T' (purple), 'Teleport' (cyan), 'Hughes' (orange), '1997' (green), 'yesterday' (orange), '1996' (green), 'WorldCom Inc' (yellow), and 'Mike Armstrong' (blue). To the right of the text is a 'Coreference data' panel. It has a 'Default' section with a list of entities, each with a checkbox and a colored box matching the text. The entities are: AT&T (checked, purple), Teleport (checked, cyan), Hughes (unchecked, orange), 1997 (unchecked, green), yesterday (unchecked, orange), 1996 (unchecked, green), WorldCom Inc (unchecked, yellow), and Mike Armstrong (checked, blue). The 'Mike Armstrong' entry is highlighted with a blue border. At the bottom of the interface are two buttons: 'Annotations Editor' and 'Features Editor'.

Messages | GATE corpus_0003D | ANNIE_0001E | example document

Text | Annotations | Annotation Sets | Coreference | Print | 🔍

"It's an indication that **Mike Armstrong** is serious about local competition and serious about getting moving," said Anna-Marie Kovacs, an analyst for the brokerage firm Janney Montgomery Scott Inc. "**AT&T** and **Teleport** are going after the business market, which is where local companies make a lot of their money."

The deal also will give three major cable television companies, which are majority owners of **Teleport**, a collective 10 percent stake in **AT&T**.

By acquiring **Teleport**, **AT&T** can offer business customers local and long-distance telephone service, and data and Internet access, under its own brand name. Using **Teleport**'s local facilities, the company also would be able to reduce the fees it pays to local phone companies for access to local telephone customers.

"It's going to permit us to be much more cost-effective as we go for that local business," **Armstrong** said at a news briefing. "This has competition and growth written all over it." **AT&T** is paying for **Teleport** with its stock. **Teleport** shareholders will receive 0.943 **AT&T** shares for each of their **Teleport** shares, putting the deal at \$ 59 a share based on **AT&T**'s closing price yesterday of \$ 62.62 1/2 a share, up \$ 2.62 1/2. **Teleport** closed down \$ 3.62 1/2 at \$ 54.12 1/2 a share. The companies expect the deal, which must be approved by regulators

Coreference data

- Default
 - AT&T**
 - Teleport**
 - Hughes**
 - 1997**
 - yesterday**
 - 1996**
 - WorldCom Inc**
 - Mike Armstrong**

Annotations Editor | Features Editor

NER—automatic approaches

- Learning of statistical models or symbolic rules
 - Use of annotated text corpus
 - Manually annotated
 - Semi-automatically annotated
- ML approaches frequently break down the NE task into two parts:
 - Recognising the entity boundaries
 - Classifying the entities in the NE categories

NER – automatic approaches

- Tokens in text are often coded with the IOB scheme
 - O – outside, B-XXX – first word in NE, I-XXX – all other words in NE

e.g.

| | |
|-----------|-------|
| Argentina | B-LOC |
| played | O |
| with | O |
| Del | B-PER |
| Bosque | I-PER |

- Probabilities:
 - Simple:
 - $P(\text{tag } i \mid \text{token } i)$
 - With external evidence:
 - $P(\text{tag } i \mid \text{token } i-1, \text{token } i, \text{token } i+1)$

NER—automatic approaches

- Decision trees
 - Tree-oriented sequence of tests in every word
 - Determine probabilities of having a IOB tag
 - Use training data
 - Viterbi, ID3, C4.5 algorithms
 - Select most probable tag sequence
 - SEKINE et al (1998)
 - BALUJA et al (1999)
 - F-measure: 90%

NER – automatic approaches

- HMM-*Generative model*
 - Markov models, Viterbi
 - Works well when large amount of data is available Nymble (1997) / IdentiFinder (1999)
- Maximum Entropy (ME)-*Discriminative model*
 - Separate, independent probabilities for every evidence (external and internal features) are merged multiplicatively
 - MENE (NYU-1998)
 - Capitalization, many lexical features, type of text
 - F-Measure: 89%

ML features

- The choice of features
 - Lexical features (words)
 - Part-of-speech
 - Orthographic information
 - Affixes (prefix and suffix of any word)
 - Gazetteers
- External, unmarked data is useful to derive gazetteers and for extracting training instances

IdentiFinder [Bikel et al 99]

- Based on Hidden Markov Models
- 7 regions of HMM—one for each *MUC type*, *not-name*, *begin-sentence* and *end-sentence*
- Features
 - Capitalisation
 - Numeric symbols
 - Punctuation marks
 - Position in the sentence
 - 14 features in total, combining above info, e.g.,
containsDigitAndDash (09-96), containsDigitAndComma
(23,000.00)

MENE [Borthwick et al 98]

- Rule-based NE + ML based NE- achieve better performance
- Tokens tagged as: XXX_start, XXX_continue, XXX_end, XXX_unique, other (non-NE), where XXX is an NE category
- Uses Maximum Entropy (ME)
 - One only needs to find the best features for the problem
 - ME estimation routine finds the best relative weights for the features

MENE (2)

- Features

- Binary features—“token begins with capitalised letter”, “token is a four-digit number”
- Lexical features—dependencies on the surrounding tokens (window ± 2) e.g., “Mr” for people, “to” for locations
- Dictionary features—equivalent to gazetteers (first names, company names, dates, abbreviations)
- External systems—whether the current token is recognised as a NE by a rule-based system

MENE (3)

- MUC-7 formal run corpus
 - MENE – *84.2%* F-measure
 - Rule-based systems – *86% - 91 %* F-measure
 - MENE + rule-based systems – *92%* F-measure
- Learning curve
 - 20 docs – 80.97% F-measure
 - 40 docs – 84.14% F-measure
 - 100 docs – 89.17% F-measure
 - 425 docs – 92.94% F-measure

Named Entity Recognition: Maximum Entropy Approach Using Global Information

(Chieu and Ng, 2003)

Global Information

- Local Context is insufficient
 - “**Mary Kay** Names Vice Chairman...”

- Global Information is useful
 - “Richard C. Bartlett was named to the newly created position of vice chairman of **Mary Kay Corp.**”

Named Entity Recognition

- Modeled as a classification problem
- Each token is assigned one of 29 ($= 7*4 + 1$) classes:
 - person_begin, person_continue, person_end,
person_unique
 - org_begin, org_continue, org_end, org_unique,
 - ...
 - nn (not-a-name)

Named Entity Recognition

Consuela Washington , a longtime

person_begin person_end nn nn nn

House staffer ... the Securities and

org_unique nn nn org_begin org_continue

Exchange Commission in the Clinton ...

org_continue org_end nn nn person_unique

Maximum Entropy Modeling

The distribution p^* in the conditional ME framework:

$$p^*(o | h) = \frac{1}{Z(h)} \prod_{j=1}^k \alpha_j^{f_j(h,o)}$$

$f_j(h,o)$: binary feature

α_j : parameter / weight of each feature

Java-based opennlp maxent package:

<http://maxent.sourceforge.net>

Checking for Valid Sequence

- To discard invalid sequences like:
 - person_begin location_end ...
- Transition probability $P(c_i | c_{i-1}) = 1$ if a valid transition, 0 otherwise
 - Dynamic programming to determine the valid sequence of classes with highest probability

$$P(c_1, \dots, c_n | s, D) = \prod_{i=1}^n P(c_i | s, D) * P(c_i | c_{i-1})$$

Local Features

- Case and zone
 - initCaps, allCaps, mixedCaps
 - TXT, HL, DATELINE, DD
- First word
- Word string
- Out-of-vocabulary
 - WordNet

Local Features

- InitCapPeriod (e.g., *Mr.*)
- OneCap (e.g., *A*)
- AllCapsPeriod (e.g., *CORP.*)
- ContainDigit (e.g., *AB3, 747*)
- TwoD (e.g., *99*)
- FourD (e.g., *1999*)
- DigitSlash (e.g., *01/01*)
- Dollar (e.g., *US\$20*)
- Percent (e.g., *20%*)
- DigitPeriod (e.g., *\$US3.20*)

Local Features

- Dictionary word lists
 - Person first names, person last names, organization names, location names
- Person prefix list (e.g., *Mr.*, *Dr.*), corporate suffix list (e.g., *Corp.*, *Inc.*)
 - Obtained from training data
- Month names, Days of the week, Numbers

Global Features

- Initcaps of other occurrences

Even Daily News have made the same mistake

They criticised **Daily News** for missing something **even** a boy would have noticed....

Global Features

- Person prefix and corporate suffix of other occurrences

Mary Kay Names Vice Chairman

Richard C. Bartlett was named to the newly created position of vice chairman of **Mary Kay Corp.**

Global Features

- Acronyms

The **Federal Communications Commission** killed

that plan last year

The company is still trying to challenge the **FCC's**
earlier decision

Global Features

- Sequence of initial caps

[HL] **First Fidelity** Unit Heads Named

[TXT] Both were executive vice presidents at **First Fidelity**.

NER – other approaches

- Hybrid systems
 - Combination of techniques
 - IBM's Intelligent Miner: Nominator + DB/2 data mining
 - WordNet hierarchies
 - MAGNINI et al. (2002)
 - Stacks of classifiers
 - Adaboost algorithm
 - Bootstrapping approaches
 - Small set of seeds
 - Memory-based ML, etc.

NER in various languages

- Arabic
 - TAGARAB (1998)
 - Pattern-matching engine + morphological analysis
 - Lots of morphological info (no differences in orthographic case)
- Bulgarian
 - OSENOVA & KOLKOVSKA (2002)
 - Handcrafted cascaded regular NE grammar
 - Pre-compiled lexicon and gazetteers
- Catalan
 - CARRERAS et al. (2003b) and MÁRQUEZ et al. (2003)
 - Extract Catalan NEs with Spanish resources (F-measure 93%)
 - Bootstrap using Catalan texts

NER in various languages

- Chinese & Japanese
 - Many works
 - Special characteristics
 - Character or word-based
 - No capitalization
 - CHINERS (2003)
 - Sports domain
 - Machine learning
 - Shallow parsing technique

NER in various languages

- ASAHARA & MATSMUTO (2003)
 - Character-based method
 - Support Vector Machine
 - 87.2% F-measure in the IREX (outperformed most word-based systems)
- Dutch
 - DE MEULDER et al. (2002)
 - Hybrid system
 - Gazetteers, grammars of names
 - Machine Learning Ripper algorithm

NER in various languages

- French
 - BÉCHET et al. (2000)
 - Decision trees
 - Le Monde news corpus
- German
 - Non-proper nouns also capitalized
 - THIELEN (1995)
 - Incremental statistical approach
 - 65% of corrected disambiguated proper names

NER in various languages

- Greek
 - KARKALETIS et al. (1998)
 - English – Greek GIE (Greek Information Extraction) project
 - GATE platform
- Italian
 - CUCCHIARELLI et al. (1998)
 - Merge rule-based and statistical approaches
 - Gazetteers
 - Context-dependent heuristics
 - ECRAN (Extraction of Content: Research at Near Market)
 - GATE architecture
 - Lack of linguistic resources: 20% of NEs undetected

NER in various languages

- Korean
 - CHUNG et al. (2003)
 - Rule-based model, Hidden Markov Model, boosting approach over unannotated data
- Portuguese
 - SOLORIO & LÓPEZ (2004, 2005)
 - Adapted CARRERAS et al. (2002b) spanish NER
 - Brazilian newspapers

NER in various languages

- Serbo-croatian
 - NENADIC & SPASIC (2000)
 - Hand-written grammar rules
 - Highly inflective language
 - Lots of lexical and lemmatization pre-processing
 - Dual alphabet (Cyrillic and Latin)
 - Pre-processing stores the text in an independent format
- Spanish
 - CARRERAS et al. (2002b)
 - Machine Learning, AdaBoost algorithm
 - BIO and OpenClose approaches

NER in various languages

- Swedish
 - SweNam system (DALIANIS & ASTROM, 2001)
 - Perl
 - Machine Learning techniques and matching rules
- Turkish
 - TUR et al (2000)
 - Hidden Markov Model and Viterbi search
 - Lexical, morphological and context clues

Named Entity Recognition

- Multilingual approaches
 - Goals - CUCERZAN & YAROWSKY (1999)
 - To handle basic language-specific evidences
 - To learn from small NE lists (about 100 names)
 - To process large and small texts
 - To have a good class-scalability (to allow the definition of different classes of entities, according to the language or to the purpose)
 - To learn incrementally, storing learned information for future use

Named Entity Recognition

- Multilingual approaches
 - GALLIPI (1996)
 - Machine Learning
 - English, Spanish, Portuguese
 - ECRAN (Extraction of Content: Research at Near Market)
 - REFLEX project (2005)
 - the US National Business Center

Named Entity Recognition

- Multilingual approaches
 - POIBEAU (2003)
 - Arabic, Chinese, English, French, German, Japanese, Finnish, Malagasy, Persian, Polish, Russian, Spanish and Swedish
 - UNICODE
 - Language independent architecture
 - Rule-based, machine-learning
 - Sharing of resources (dictionary, grammar rules...) for some languages
 - BOAS II (2004)
 - University of Maryland Baltimore County
 - Web-based
 - Pattern-matching
 - No large corpora

NER – other topics

- Character vs. word-based
 - JING et al. (2003)
 - Hidden Markov Model classifier
 - Character-based model better than word-based model
- NER translation
 - Cross-language Information Retrieval (CLIR), Machine Translation (MT) and Question Answering (QA)
- NER in speech
 - No punctuation, no capitalization
 - KIM & WOODLAND (2000)
 - Up to 88.58% F-measure
- NER in Web pages
 - wrappers

NER in Indian Languages

Problems for NER in Indian Languages

- Lacks **capitalization information**
- More **diverse Indian person names**
 - Lot of person names appear in the **dictionary** with other **specific meanings**
 - For e.g., *KabiTA* (**Person name** vs. **Common noun** with meaning ‘**poem**’)
- High **inflectional nature** of Indian languages
 - Richest and most challenging sets of **linguistic** and **statistical features** resulting in **long and complex wordforms**
- **Free word order** nature of the Indian languages
- **Resource-constrained environment** of Indian languages
 - PoS taggers, morphological analyzers, name lists etc. are not available in the web
- **Non-availability** of sufficient published works

NER in Indian Languages

- LI and McCallum (2004)-Hindi
 - CRF model using feature induction technique to automatically construct the features
 - Features:
 - Word text, character n-grams ($n=2, 3, 4$), word prefix and suffix of lengths 2,3,4
 - 24 Hindi gazetteer lists
 - Features at the current, previous and next sequence positions were made available
 - Dataset: 601 BBC and 27 EMI Hindi documents
 - Performance
 - *F-measure* of 71.5% with an early stopping point of 240 iterations of L-BFGS for the 10-fold cross validation experiments

NER in Indian Languages

- Saha et al. (2008)-Hindi
 - ME model
 - Features:
 - Statistical and linguistic feature sets
 - Hindi gazetteer lists
 - Semi-automatic induction of context patterns
 - Context patterns as features of the MaxEnt method
 - Dataset: 243K words of Dainik Jagaran (training)
25K (test)
 - Performance
 - *F-measure* of 81.52%

NER in Indian Languages

- Patel et al. (2008)-Hindi and Marathi
 - Inductive Logic Programming (ILP) based techniques for automatically extracting rules for NER from tagged corpora and background knowledge
 - Dataset: 54340 (Marathi), 547138 (Hindi)
 - Performance
 - *PER: 67%, LOC: 71% and ORG: 53% (Marathi)*
 - *PER: 82%, LOC: 48% and ORG: 55% (Hindi)*
 - Advantages over rule-based system
 - *development time reduces by a factor of 120 compared to a linguist doing the entire rule development*
 - *a complete and consistent view of all significant patterns in the data at the level of abstraction*

NER in Indian Languages

- Ekbal and Saha (2011)-Bengali, Hindi, Telugu and Oriya
 - Genetic algorithm based weighted ensemble
 - Classifiers: ME, CRF and SVM
 - Features:
 - Word text, word prefix and suffix of lengths 1,2,3; PoS
 - Context information, various orthographic features etc.
 - Dataset: Bengali (Training: 312,947; Test: 37,053)
Hindi (Training: 444,231; Test: 58,682)
Telugu (Training: 57,179; Test: 4,470)
Oriya (Training: 93,573; Test: 2,183)
 - Performance
 - *F-measures: Bengali (92.15%), Hindi (92.20%), Telugu (84.59%) and Oriya (89.26%)*

NER in Indian Languages

- Ekbal and Saha (2012)-Bengali, Hindi and Telugu
 - Multiobjective Genetic algorithm based weighted ensemble
 - Classifiers: ME, CRF and SVM
 - Features:
 - Word text, word prefix and suffix of lengths 1,2,3; PoS
 - Context information, various orthographic features etc.
 - Dataset: Bengali (Training: 312,947; Test: 37,053)
Hindi (Training: 444,231; Test: 58,682)
Telugu (Training: 57,179; Test: 4,470)
Oriya (Training: 93,573; Test: 2,183)
 - Performance
 - *F-measures: Bengali (92.46%), Hindi (93.20%), Telugu (86.54%)*

NER in Indian Languages

- Shishtla et al. (2008)- Telugu and Hindi
 - CRF
 - Character-n gram approach is more effective than word-based model
 - Features
 - Word-internal features, PoS, chunk etc.
 - No external resources
 - Datasets: Telugu (45,714 tokens); Hindi ((45,380 tokens)
 - Performance
 - F-measures: Telugu (49.62%), Hindi (45.07%)

NER in Indian Languages

- Vijayakrishna and Sobha (2008)
 - CRF
 - Tourism domain with 106 hierarchical tags
 - Features
 - Roots of words, PoS, dictionary of NEs, patterns of certain types of NEs (date, time, money etc.) etc
 - Performance
 - 80.44%

NER in Indian Languages

- Saha et al. (2008)- Hindi
 - Maximum Entropy
 - Features
 - Statistical and linguistics features
 - Word clustering
 - Clustering used for feature reduction in Maximum Entropy
- -Datasets: 243K Hindi newspaper “Dainik Jagaran”.
 - Performance
 - F-measures: 79.03% (approximately 7% improvement with Clusters)

Other works in Indian Languages NER

- Gali et al. (2008)-Bengali, Hindi, Telugu and Oriya
 - CRF
- Kumar and Kiran (2008)-Bengali, Hindi, Telugu and Oriya
 - CRF
- Srikanth and Murthy (2008) –Telugu
 - CRF
- Goyal (2008)-Hindi
 - CRF
- Nayan et al. (2008)-Hindi
 - Phonetic matching technique

Other works in Indian Languages NER

- Ekbal et al. (2008)-Bengali
 - CRF
- Saha et al. (2009)-Hindi
 - Semi-supervised approach
- Saha et al. (2010)-Hindi
 - SVM with string based kernel function
- Ekbal and Saha (2010)-Bengali, Hindi and Telugu
 - GA based classifier ensemble selection
- Ekbal and Saha (2011)-Bengali, Hindi and Telugu
 - Multiobjective simulated annealing approach for classifier ensemble

Other works in Indian Languages NER

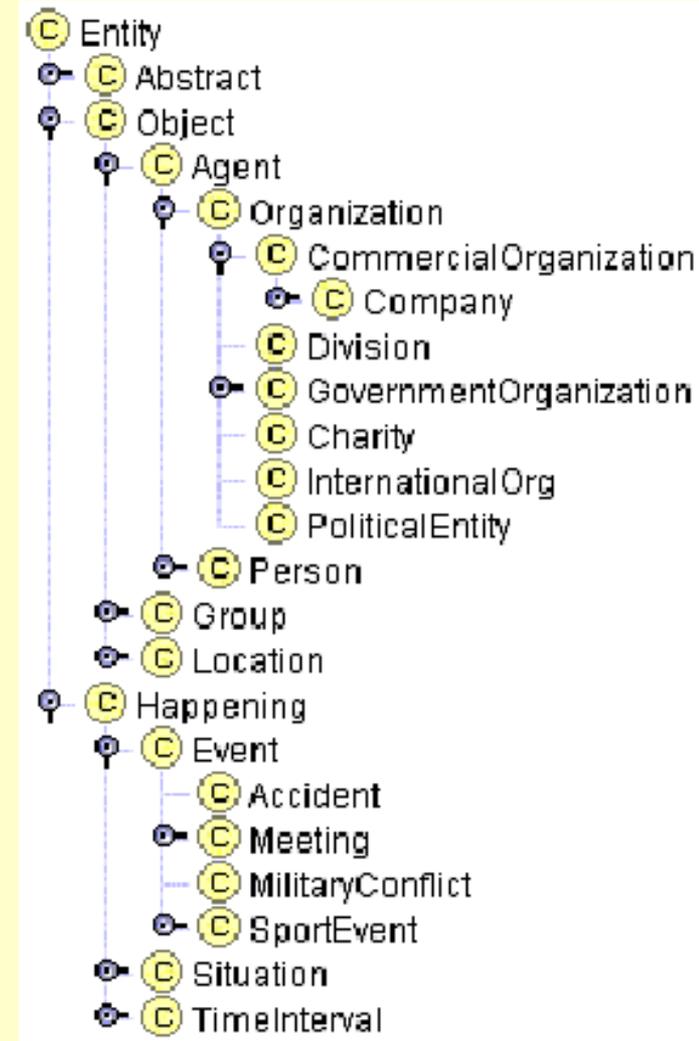
- Saha et al. (2012)-Hindi and Bengali
 - Comparative techniques for feature reductions
- Ekbal and Saha (2012)-Bengali, Hindi and Telugu
 - Multiobjective approach for feature selection and classifier ensemble
- Ekbal et al. (2012)-Hindi and Bengali
 - Active learning
 - Effective in a resource-constrained environment

Shared Tasks on Indian Language NER

- ICON 2013 Tool Contest
(<http://ltrc.iiit.ac.in/icon/2013/nlptools/>)
- NERSSEAL Shared Task 2008-
(<http://ltrc.iiit.ac.in/ner-ssea-08/index.cgi?topic=2>)
- NLP AI ML Contest 2007-
(http://ltrc.iiit.ac.in/nlpai_contest07/cgi-bin/index.cgi)

Evaluating Richer NE Tagging

- Need for new metrics when evaluating hierarchy/ontology-based NE tagging
- Need to take into account distance in the hierarchy
- Tagging a company as a charity is less wrong than tagging it as a person



Ensemble learning: An introduction

Drawbacks of Single Classifier

- The “best” classifier not necessarily the ideal choice
- For solving a classification problem, many individual classifiers with different parameters are trained
 - The “best” classifier will be selected according to some criteria e.g., *training accuracy* or *complexity of the classifiers*
- Problems: Which one is the best?
 - Maybe more than one classifiers meet the criteria (e.g. same training accuracy), especially in the following situations:
 - Without sufficient training data
 - Learning algorithm leads to different local optima easily

Drawbacks of Single Classifier

- Potentially valuable information may be lost by discarding the results of less-successful classifiers

E.g., the discarded classifiers may correctly classify some samples

- Other drawbacks
 - Final decision must be wrong if the output of selected classifier is wrong
 - Trained classifier may not be complex enough to handle the problem

Ensemble Learning

- Employ multiple learners and combine their predictions
- Methods of combination:
 - Bagging, boosting, voting
 - Error-correcting output codes
 - Stacked generalization
 - Cascading
 - ...
- **Advantage:** improvement in predictive accuracy
- **Disadvantage:** it is difficult to understand an ensemble of classifiers

Why Do Ensembles Work?

Dietterich(2002) showed that ensembles overcome three problems:

- ***Statistical Problem-*** arises when the hypothesis space is too large for the amount of available data. Hence, there are many hypotheses with the same accuracy on the data and the learning algorithm chooses only one of them! There is a risk that the accuracy of the chosen hypothesis is low on unseen data!
- ***Computational Problem-*** arises when the learning algorithm cannot guarantee finding the best hypothesis.
- ***Representational Problem-*** arises when the hypothesis space does not contain any good approximation of the target class(es).

T.G. Dietterich, Ensemble Learning, 2002

Categories of Ensemble Learning

- Methods for Independently Constructing Ensembles
 - Bagging
 - Randomness Injection
 - Feature-Selection Ensembles
 - Error-Correcting Output Coding
- Methods for Coordinated Construction of Ensembles
 - Boosting
 - Stacking
 - Co-training

Some Practical Advices

- If the classifier is **unstable** (i.e, decision trees) then apply bagging!
- If the classifier is **stable and simple** (e.g. Naïve Bayes) then apply boosting!
- If the classifier is **stable and complex** (e.g. Neural Network) then apply randomization injection!
- If you have many classes and a binary classifier then try error-correcting codes! If it does not work then use a complex binary classifier!

*Evolutionary Algorithms for Classifier
Ensemble*

Evolutionary Algorithms in NLP

- Good Review (L. Araujo, 2007)
- Natural language tagging- Alba, G. Luque, and L. Araujo (2006)
- Grammar Induction-T. C. Smith and I. H. Witten (1995)
- Phrase-structure-rule of natural language-W. Wang and Y. Zhang (2007)
- Information retrieval-R. M. Losee (2000)
- Morphology -D. Kazakov (1997)
- Dialogue systems-D. Kazakov (1998)
- Grammar inference -M. M. Lankhors (1994)
- Memory-based language processing (A. Kool, W. Daelemans, and J. Zavrel., 2000)

Evolutionary Algorithms in NLP

- Anaphora resolution: Veronique Hoste (2005), Ekbal et al. (2011), Saha et al. (2012)
- Part-of-Speech tagging: Araujo L (2002)
- Parsing: Araujo L (2004)
- Document clustering: Casillas A et al. (2003)
- Summarization: Andersson L (2004)
- Machine Translation : Jun Suzuki (2012)
- NER: Ekbal and Saha (2010; 2011; 2012 etc.)

Genetic Algorithm: Quick Overview

- Randomized search and optimization technique
- Evolution produces good individuals, similar principles might work for solving complex problems
- Developed: USA in the 1970's by J. Holland
- Got popular in the late 1980's
- Early names: J. Holland, K. DeJong, D. Goldberg
- Based on ideas from *Darwinian Evolution*
- Can be used to solve a variety of problems that are not easy to solve using other techniques

GA: Quick Overview

- Typically applied to:
 - discrete optimization
- Attributed features:
 - not too fast
 - good heuristic for combinatorial problems
- Special Features:
 - Traditionally emphasizes combining information from good parents (*crossover*)
 - many variants, e.g., reproduction models, operators

How is GA different from other traditional techniques?

- GAs work with a coding of parameter set, not the parameters themselves
- GAs search from a population of points, not a single point
- GAs use payoff (objective function) information, not derivatives or other auxiliary knowledge
- GAs use probabilistic transition rules, not deterministic rules

Genetic Algorithm: Similarity with Nature

| | | |
|---|---|------------------------|
| Genetic Algorithms | ↔ | Nature |
| A solution (phenotype) | | Individual |
| Representation of a solution (<i>genotype</i>) | | Chromosome |
| Components of the solution | | Genes |
| Set of solutions | | Population |
| Survival of the fittest (<i>Selection</i>) | | Darwins theory |
| Search operators | | Crossover and mutation |
| Iterative procedure | | Generations |

Basic Steps of Genetic Algorithm

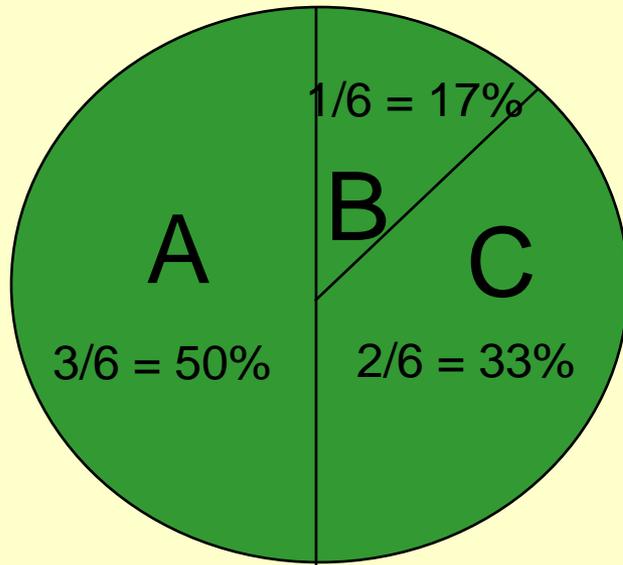
1. $t = 0$
 2. initialize population $P(t)$ /* $Popsiz e = |P|$ */
 3. for $i = 1$ to $Popsiz e$
 compute fitness $P(t)$
 4. $t = t + 1$
 5. if termination criterion achieved go to step 10
 6. select (P)
 7. crossover (P)
 8. mutate (P)
 9. go to step 3
 10. output best chromosome and stop
- End

Example population

| No. | Chromosome | Fitness |
|-----|------------|---------|
| 1 | 1010011010 | 1 |
| 2 | 1111100001 | 2 |
| 3 | 1011001100 | 3 |
| 4 | 1010000000 | 1 |
| 5 | 0000010000 | 3 |
| 6 | 1001011111 | 5 |
| 7 | 0101010101 | 1 |
| 8 | 1011100111 | 2 |

GA operators: Selection

- Main idea: better individuals get higher chance
 - Chances proportional to fitness
 - Implementation: roulette wheel technique
 - » Assign to each individual a part of the roulette wheel
 - » Spin the wheel n times to select n individuals



fitness(A) = 3

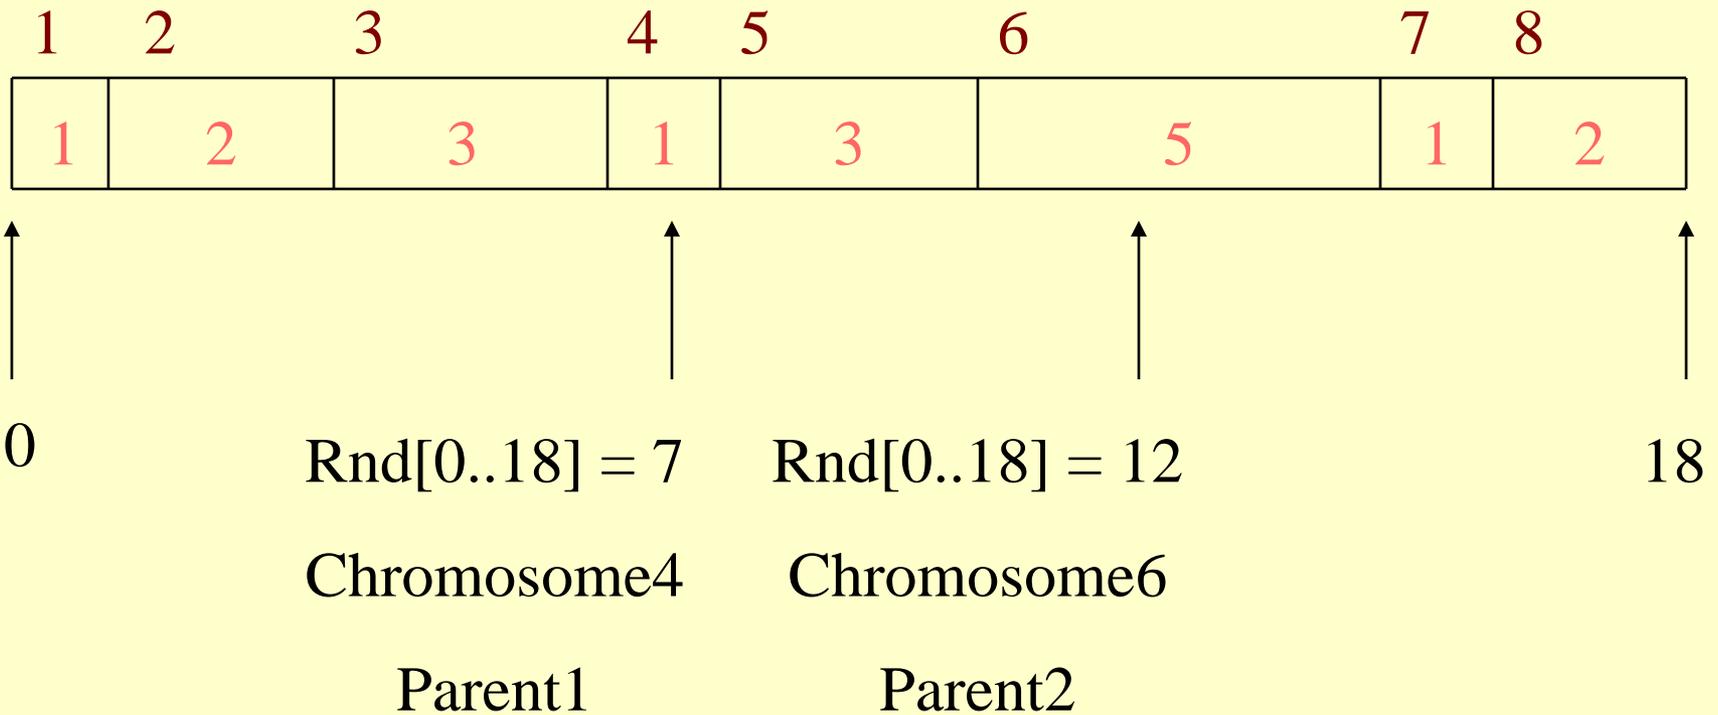
fitness(B) = 1

fitness(C) = 2

GA operator: Selection

- Add up the fitness's of all chromosomes
- Generate a random number R in that range
- Select the first chromosome in the population that -when all previous fitness's are added including the current one- gives you at least the value R

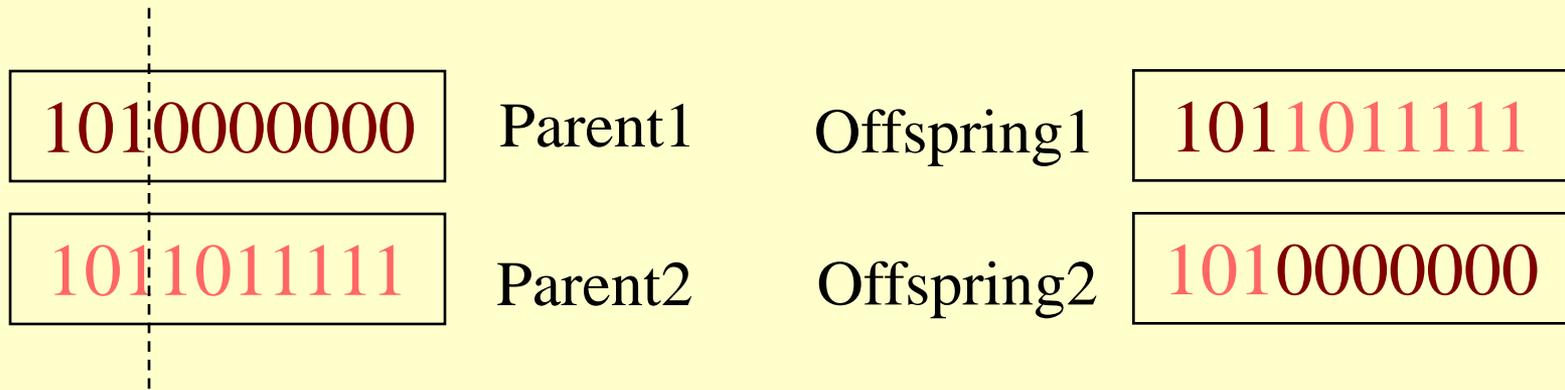
Roulette Wheel Selection



GA operator: Crossover

- Choose a random point on the two parents
- Split parents at this crossover point
- With some high probability (*crossover rate*) apply crossover to the parents
 - P_c typically in range (0.6, 0.9)
- Create children by exchanging tails

Crossover - Recombination

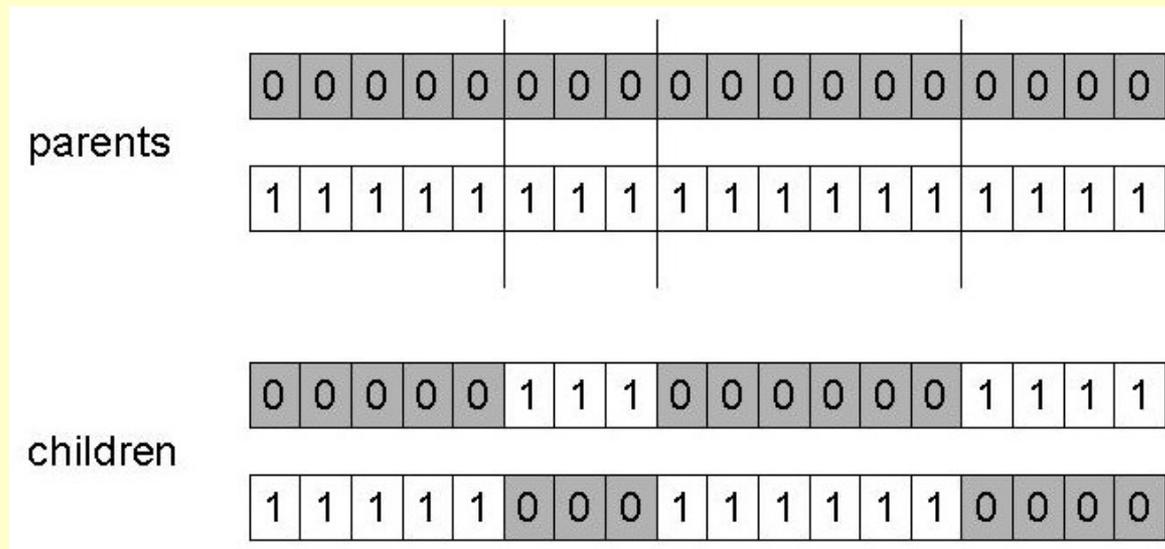


Crossover
single point -
random

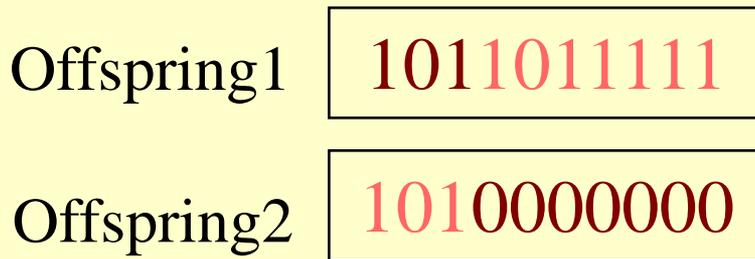
Single Point Crossover

n-point crossover

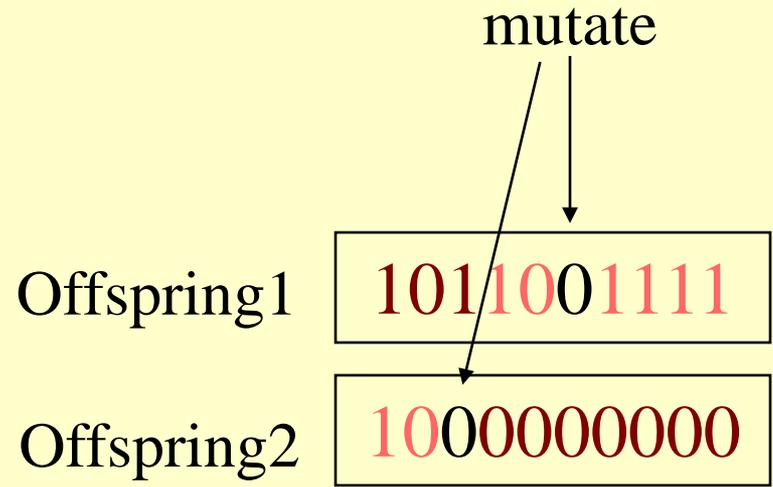
- Choose n random crossover points
- Split along those points
- Glue parts, alternating between parents
- Generalisation of 1 point (still some positional bias)



Mutation



Original offspring



Mutated offspring

With some small probability (the *mutation rate*) flip each bit in the offspring (*typical values between 0.1 and 0.001*)

A. Ekbal and S. Saha (2011). Weighted Vote-Based Classifier Ensemble for Named Entity Recognition: A Genetic Algorithm-Based Approach. ACM Transactions on Asian Language Information Processing (ACM TALIP), Vol. 2(9),

DOI=10.1145/1967293.1967296

<http://doi.acm.org/10.1145/1967293.1967296>

Weighted Vote based Classifier Ensemble

- **Motivation**
 - All classifiers are not equally good to identify all classes
- **Weighted voting:** Weights of voting vary among the classes for each classifier
 - *High*: Classes for which the classifier perform good
 - *Low*: Classes for which it's output is not very reliable
- **Crucial issue:** Selection of appropriate weights of votes per classifier

Problem Formulation

Let *no. of classifiers*= N , and *no. of classes*= M

Find the weights of votes V per classifier optimizing a function $F(V)$

- V : an real array of size $N \times M$

- $V(i, j)$: weight of vote of the i th classifier for the j th class

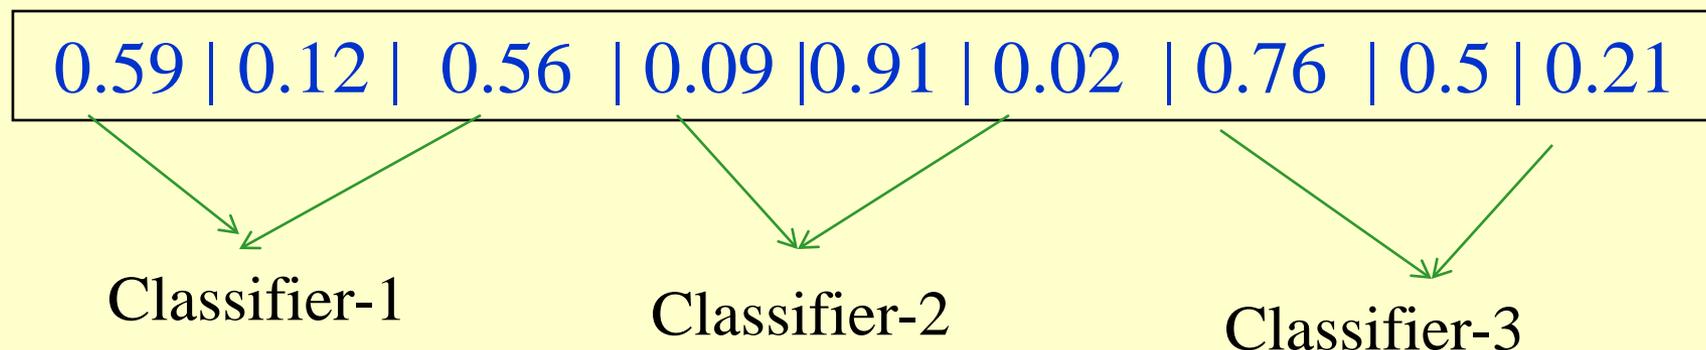
- $V(i, j) \in [0, 1]$ denotes the degree of confidence of the i th classifier for the j th class

maximize $F(B)$;

$F \in \{recall, precision, F\text{-measure}\}$ and B is a subset of A

Here, $F1 = F\text{-measure}$

Chromosome representation



- Real encoding used
- Entries of chromosome randomly initialized to a real (r) between 0 and 1: $r = \text{rand}() / \text{RAND_MAX} + 1$
- If the population size P then all the P number of chromosomes of this population are initialized in the above way

Fitness Computation

Step-1: For M classifiers, F_i $i = 1$ to M be the F-measure values

Step-2: Train each classifier with $2/3$ training data and test with the remaining $1/3$ part

Step-3: For ensemble output of the $1/3$ test data, apply weighted voting to the outputs of M classifiers

(a). Weight of the output label provided by the i th classifier = $I(m, i)$

Here, $I(m, i)$ is the entry of the chromosome corresponding to m th classifier and i th class

(b). Combined weight of a class for a word w

$$f(c_i) = \sum I(m, i) \times F_m, \quad \forall m = 1 \text{ to } M \text{ and } op(w, m) = c_i$$

Fitness Computation

$Op(w, m)$: output class produced by the m th classifier for word w

Class receiving the maximum weight selected as the joint decision

Step-4: Compute the overall F-measure value for 1 / 3 data

Step-5: Steps 3 and 4 repeated to perform 3-fold cross validation

Step-6: Objective function or fitness function = $F\text{-measure}_{avg}$

Objective: Maximize the objective function using search capability of GA

Other Parameters

- **Selection**
 - Roulette wheel selection (*Holland, 1975; Goldberg, 1989*)
- **Crossover**
 - Normal Single-point crossover (Holland, 1975)
- **Mutation**
 - Probability selected adaptively (*Srinivas and Patnaik, 1994*)
 - Helps GA to come out from local optimum

Termination Condition

- Execute the processes of *fitness computation*, *selection*, *crossover*, and *mutation* for a maximum number of generations
- *Best solution*-Best string seen up to the last generation
- Best solution indicates
 - Optimal voting weights for all classes in each classifier
- Elitism implemented at each generation
 - Preserve the best string seen up to that generation in a location outside the population
 - Contains the most suitable classifier ensemble

NE Features

- **Context Word:** Preceding and succeeding words
- **Word Suffix**
 - Not necessarily **linguistic suffixes**
 - **Fixed length** character strings stripped from the **endings of words**
 - **Variable length** suffix -binary valued feature
- **Word Prefix**
 - **Fixed length** character strings stripped from the **beginning of the words**
- **Named Entity Information:** **Dynamic NE tag (s)** of the **previous word (s)**

NE Features

- **First Word (binary valued feature)**: Check whether the current token is the first word in the sentence
- **Length (binary valued)**: Check whether the length of the current word less than **three** or not (shorter words rarely NEs)
- **Position (binary valued)**: Position of the word in the sentence
- **Infrequent (binary valued)**: Infrequent words in the training corpus most probably NEs

NE Features

- Digit features: Binary-valued
 - Presence and/or the exact number of digits in a token
 - **CntDgt** : Token contains digits
 - **FourDgt**: Token consists of four digits
 - **TwoDgt**: Token consists of two digits
 - **CnsDgt**: Token consists of digits only
- Combination of digits and punctuation symbols
 - **CntDgtCma**: Token consists of digits and comma
 - **CntDgtPrd**: Token consists of digits and periods

NE Features

- Combination of digits and symbols
 - **CntDgtSlsh**: Token consists of digit and slash
 - **CntDgtHph**: Token consists of digits and hyphen
 - **CntDgtPrctg**: Token consists of digits and percentages
- Combination of digit and special symbols
 - **CntDgtSpl**: Token consists of digit and special symbol such as \$, # etc.

NE Features

- **Part of Speech (POS) Information:** POS tag(s) of the current and/or the surrounding word(s)
 - **SVM-based POS tagger** (Ekbal and Bandyopadhyay, 2008)
 - **Coarse-grained POS tagger**
 - **Nominal, PREP** (Postpositions) and **Other**
- **Gazetteer based features (binary valued):** Several features extracted from the **gazetteers**
- **Content words in surrounding contexts**-*Exploits global context information*

Datasets

- Web-based Bengali news Corpus (Ekbal and Bandyopadhyay, 2008, *Language Resources and Evaluation of Springer*)
 - *34 million* wordforms
 - News data collection of 5 years
- NE annotated corpus for Bengali
 - Manually annotated 250K wordforms
 - IJCNLP-08 Shared Task on NER for South and South East Asian Languages (available at <http://ltrc.iiit.ac.in/ner-ssea-08>)
- NE annotated datasets for Hindi and Telugu
 - NERSSEAL shared task

NE Tagset

- Reference Point- CoNLL 2003 shared task tagset
- Tagset: 4 NE tags
 - Person name
 - Location name
 - Organization name
 - Miscellaneous name (*date, time, number, percentages, monetary expressions and measurement expressions*)
- IJCNLP-08 NERSSEAL Shared Task Tagset: Fine-grained 12 NE tags (available at <http://ltrc.iiit.ac.in/ner-ssea-08>)
- Tagset Mapping (12 NE tags → 4 NE tags)
 - ☐ NEP → Person name
 - ☐ NEL → Location name
 - ☐ NEO → Organization name
 - ☐ NEN [number], NEM [Measurement] and NETI [time] → Miscellaneous name
 - ☐ NETO [title-object], NETE [term expression], NED [designations], NEA [abbreviations], NEB [brand names], NETP [title persons]

Training and Test Datasets

| Language | #Words in training | #NEs in training | #Words in test | #NEs in test |
|----------|--------------------|------------------|----------------|--------------|
| Bengali | 312,947 | 37,009 | 37,053 | 4,413 |
| Hindi | 444,231 | 26,432 | 32,796 | 58,682 |
| Telugu | 57,179 | 4,470 | 6,847 | 662 |
| Oriya | 93,573 | 4,477 | 2,183 | 206 |

Experiments

- Classifiers used
 - Maximum Entropy (ME): Java based OpenNLP package (<http://maxent.sourceforge.net/>)
 - Conditional Random Field: C++ based CRF++ package (<http://crfpp.sourceforge.net/>)
 - Support Vector Machine:
 - YamCha toolkit
(<http://chasen-org/taku/software/yamcha/>)
 - TinySVM-0.07
(<http://cl.aist-nara.ac.jp/taku-ku/software/TinySVM>)
 - Polynomial kernel function

Experiments

- **GA**: population size=50, number of generations=40, mutation and crossover probabilities are selected adaptively.
- **Baselines**
 - Baseline 1: Majority voting of all classifiers
 - Baseline 2: Weighted voting of all classifiers (*weight*: overall average F-measure value)
 - Baseline 3: Weighted voting of all classifiers (*weight*: F-measure value of the individual class)

Results (*Bengali*)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best Individual Classifier | 89.42 | 90.55 | 89.98 |
| Baseline-1 | 84.83 | 85.90 | 85.36 |
| Baseline-2 | 85.25 | 86.97 | 86.97 |
| Baseline-3 | 86.97 | 87.34 | 87.15 |
| Stacking | 90.17 | 91.74 | 90.95 |
| ECOC | 89.78 | 90.89 | 90.33 |
| QBC | 90.01 | 91.09 | 90.55 |
| GA based ensemble | 92.08 | 92.22 | 92.15 |

Results (*Hindi*)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best Individual Classifier | 88.72 | 90.10 | 89.40 |
| Baseline-1 | 63.32 | 90.99 | 74.69 |
| Baseline-2 | 74.67 | 94.73 | 83.64 |
| Baseline-3 | 75.52 | 96.13 | 84.59 |
| Stacking | 89.80 | 90.61 | 90.20 |
| ECOC | 90.16 | 91.11 | 90.63 |
| GA based ensemble | 96.07 | 88.63 | 92.20 |

Results (*Telugu*)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best Individual Classifier | 77.42 | 77.99 | 77.70 |
| Baseline-1 | 60.12 | 87.39 | 71.23 |
| Baseline-2 | 71.87 | 92.33 | 80.33 |
| Baseline-3 | 72.22 | 93.10 | 81.34 |
| Stacking | 77.65 | 84.12 | 80.76 |
| ECOC | 77.96 | 85.12 | 81.38 |
| GA based ensemble | 78.82 | 91.26 | 84.59 |

Results (*Oriya*)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best Individual Classifier | 86.55 | 88.03 | 87.29 |
| Baseline-1 | 86.95 | 88.33 | 87.63 |
| Baseline-2 | 87.12 | 88.50 | 87.80 |
| Baseline-3 | 87.62 | 89.12 | 88.36 |
| Stacking | 87.90 | 89.53 | 88.71 |
| ECOC | 87.04 | 88.56 | 87.79 |
| GA based ensemble | 88.56 | 89.98 | 89.26 |

Results (*English*)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best Individual Classifier | 86.16 | 85.24 | 86.31 |
| Baseline-1 | 85.75 | 86.12 | 85.93 |
| Baseline-2 | 86.20 | 87.02 | 86.61 |
| Baseline-3 | 86.65 | 87.25 | 86.95 |
| Stacking | 85.93 | 86.45 | 86.18 |
| ECOC | 86.12 | 85.34 | 85.72 |
| GA based ensemble | 88.72 | 88.64 | 88.68 |

Multiobjective Optimization

*(Simultaneous optimization of more than
one objective)*

Single vs. Multi-objective

Single Objective Optimization:

When an optimization problem involves only one objective function, the task of finding the optimal solution is called single-objective optimization

Example: Find out a **CAR** for me with Minimum cost.

Multi-objective Optimization: When an optimization problem involves more than one objective function, the task of finding one or more optimal solutions is known as multi-objective optimization.

Example: Find out a **CAR** with minimum cost and maximum comfort.

Multiobjective Optimization:

Example of purchasing a car

- *Optimizing criteria*
 - minimizing the cost, insurance premium and weight and
 - maximizing the feel good factor while in the car
- *Constraints*
 - car should have good stereo system, seats for 6 adults and a mileage of 20 kmpl
- *Decision variables*
 - the available cars
- In many real world problems we have to **simultaneously** optimize two or more different objectives which are often **competitive** in nature
 - finding a single solution in these cases is very difficult
 - optimizing each criterion separately may lead to good value of one objective while some unacceptably low value of the other objective(s)

Multiobjective Optimization: Mathematical Definition

The multiobjective optimization can be formally stated as:

Find the vector of decision variables

$$\mathbf{x} = [x_1, x_2, \dots, x_n]^T$$

which will satisfy the m inequality constraints:

$$g_i(\mathbf{x}) \geq 0, \quad i=1, 2, \dots, m,$$

And the p equality constraints

$$h_i(\mathbf{x}) = 0, \quad i=1, 2, \dots, p.$$

And simultaneously optimizes M objective functions

$$f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_M(\mathbf{x}).$$

Pareto Optimum: Definition

- A candidate is *Pareto optimal* iff:
 - It is at least as good as all other candidates for all objectives, and
 - It is better than all other candidates for at least one objective
- We would say that this candidate *dominates* all other candidates

Dominance: Definition

Given the vector of objective functions $\vec{f}(\vec{x}) = (f_1(\vec{x}), \dots, f_k(\vec{x}))$

we say that candidate \vec{x}_1 dominates \vec{x}_2 , (i.e. $\vec{x}_1 \preceq \vec{x}_2$) if:

$$f_i(\vec{x}_1) \leq f_i(\vec{x}_2) \quad \forall i \in \{1, \dots, k\}$$

and

$$\exists i \in \{1, \dots, k\} : f_i(\vec{x}_1) < f_i(\vec{x}_2)$$

(assuming we are trying to minimize the objective functions).

Pareto Optimal Set

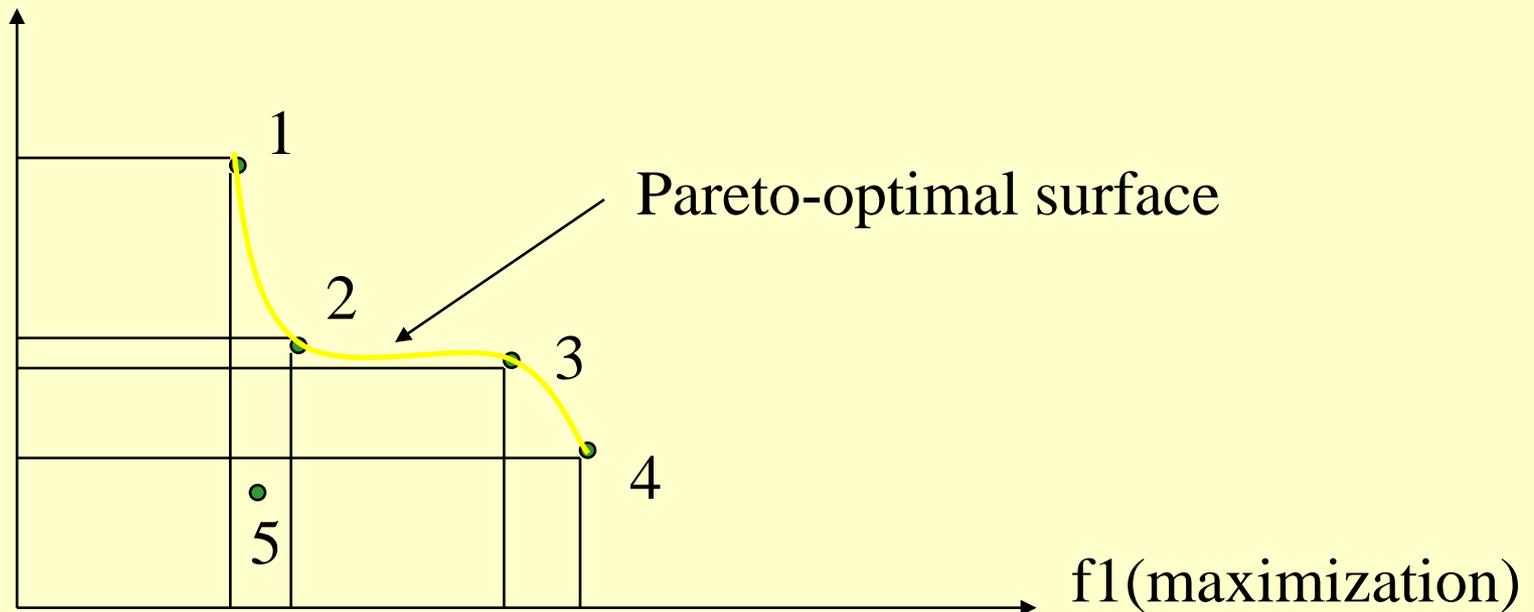
The Pareto optimal set P contains all candidates that are non-dominated. That is:

$$P := \left\{ x \in F \mid \left[\neg \exists x' \in F \right] \ni \left(\vec{f}(x') \preceq \vec{f}(x) \right) \right\}$$

where F is the set of feasible candidate solutions

Example of Dominance and Pareto-Optimality

f2(maximization)



- Here solutions 1, 2, 3 and 4 are non-dominating to each other.
- 5 is dominated by 2, 3 and 4, not by 1.

Types of Solutions

- **Non-dominated solutions**
 - Solutions that lie along the line
- **Dominated solutions**
 - Solutions that lie inside the line because there is always another solution on the line that has at least one objective that is better

Pareto-optimal Solutions

- Line is called the *Pareto front* and solutions on it are called *Pareto-optimal*
- All Pareto-optimal solutions are non-dominated

Thus, it is important in MOO to find the solutions as close as possible to the Pareto front & as far along it as possible

Why Evolutionary Algorithm for MOO?

- Evolutionary algorithms seem particularly suitable to solve MOO problems, because
 - simultaneously deal with a set of possible solutions (*population based nature*)
- Allows to find several members of the Pareto optimal set in a single run of the algorithm
 - In traditional mathematical programming techniques, we need to perform a series of separate runs
- EAs are less susceptible to the shape or continuity of the Pareto front (e.g., they can easily deal with discontinuous or concave Pareto fronts)
 - Real concerns for mathematical programming techniques

A. Ekbal and S. Saha (2012). Multiobjective Optimization for Classifier Ensemble and Feature Selection: An Application to Named Entity Recognition. International Journal on Document Analysis and Recognition (IJDAR), Vol. 15(2), 143-166, Springer

Why MOO in Classifier Ensemble?

- Single objective optimization technique : optimizes a single quality measure
 - recall, precision or F-measure at a time
- A single measure cannot capture the quality of a good ensemble reliably
- A good classifier ensemble should have it's all the parameters optimized simultaneously
- **Advantages of MOO**
 - MOO to simultaneously optimize more than one classification quality measures
 - Provides user a set of alternative solutions

Formulation of Classifier Ensemble Selection Problem

Classifier ensemble selection problem:

A: Set of N classifiers

Find a set of classifiers B that maximizes

$[F1(B), F2(B)]$

where

$F1, F2 \in \{\text{recall, precision, F-measure}\}$ and

$F1 \neq F2$

Here, $B \subseteq A$

$F1 = \text{recall}$ and $F2 = \text{precision}$

Classifier Ensemble Selection: Proposed Approach

Chromosome representation

010110111110011111

Total number of available classifiers: M

0 at position i - i th classifier does not participate in ensemble

1 at position i - i th classifier participates in ensemble

Fitness Computation

Step-1: For M classifiers, F_i $i= 1$ to M be the F-measure values

Step-2: Train each classifier with $2/3$ training data and test with the remaining $1/3$ part.

Step-3: For ensemble output of the $1/3$ test data

- a. Appropriate class is determined from the weighted voting
- b. weight = F-measure value of the respective classifier

Step-4: Calculate the overall *recall*, *precision* and *F-measure* values for $1/3$ data

Steps 2 -4 are repeated 3 times to perform 3-fold cross validation.

Step-5: Average *recall* and *precision* values are considered as two objective functions

Other Operators

- Steps of non-dominated sorting genetic algorithm (NSGA-II) are executed (Deb K et al., 2002)
- Crowded binary tournament selection
- Conventional *crossover* and *mutation*
- *Elitism*-non-dominated solutions among the parent and child populations are propagated to the next generation (Deb K, 2001)
- *Near-Pareto-optimal* strings of the last generation provide the different solutions to the ensemble problem

Selecting Solution from Pareto Optimal Front

- In MOO, the algorithms produce a *large number of non-dominated solutions* on the final *Pareto optimal front*
- Each of these solutions provides a classifier ensemble
- All the *solutions are equally important* from the algorithmic point of view
- User may want only a single solution

Selecting Solution from Pareto Optimal Front

- For every solution on the final Pareto optimal front
 - calculate the overall average *F-measure value of the classifier ensemble* for the three-fold cross-validation
- Select the solution with the maximum *F-measure value* as the best solution
- Evaluate the classifier ensemble corresponding to the best solution on the test data

Experiments

- **Classifiers used**

- Maximum Entropy (ME): Java based OpenNLP package (<http://maxent.sourceforge.net/>)
- Conditional Random Field: C++ based CRF++ package (<http://crfpp.sourceforge.net/>)
- Support Vector Machine:
 - YamCha toolkit (<http://chasen-org/taku/software/yamcha/>)
 - TinySVM-0.07 (<http://cl.aist-nara.ac.jp/taku-ku/software/TinySVM>)
 - Polynomial kernel function

Experiments

- *NSGA-II* (<http://www.iitk.ac.in/kangal/codes.shtml>): population size = 100, number of generations = 50, probability of mutation = 0.2, and probability of crossover = 0.9.
- *Baselines*
 - *Baseline 1*: Majority voting of all classifiers
 - *Baseline 2*: Weighted voting of all classifiers (weight: overall average F-measure value)

Results (*Bengali*)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best Individual Classifier | 89.42 | 90.55 | 89.98 |
| Baseline-1 | 84.83 | 85.90 | 85.36 |
| Baseline-2 | 85.25 | 86.97 | 86.97 |
| GA based ensemble | 91.08 | 91.22 | 91.15 |
| MOO based ensemble | 92.21 | 92.72 | 92.46 |

Results (*Hindi*)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best Individual Classifier | 88.72 | 90.10 | 89.40 |
| Baseline-1 | 63.32 | 90.99 | 74.69 |
| Baseline-2 | 74.67 | 94.73 | 83.64 |
| GA based ensemble | 89.92 | 91.16 | 90.54 |
| MOO based ensemble | 97.07 | 89.63 | 93.20 |

Results (*Telugu*)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best Individual Classifier | 77.42 | 77.99 | 77.70 |
| Baseline-1 | 60.12 | 87.39 | 71.23 |
| Baseline-2 | 71.87 | 92.33 | 80.83 |
| GA based ensemble | 78.02 | 90.26 | 83.69 |
| MOO based ensemble | 80.79 | 93.18 | 86.54 |

Study Materials

- *Named Entities: Recognition, Classification and Use, Special Issue of Lingvisticae Investigationes Journal*, Satoshi Sekine and Elisabete Ranchhod (Eds.), Vol. 30:1 (2007), John Benjamins Publishing Company
- Journals-CL, ACM TALIP, JAIR etc.
- All relevant conferences- ACL, COLING, EACL, IJCNLP, CiCLing , AAI, ECAI etc.
- Named Entities Workshop (NEWS)

Current Trends in NE Research

- Development of domain-independent and language-independent systems
 - Can be easily portable to different domains and languages
- Fine-grained NE classification
 - May be at the hierarchy of WordNet
 - Beneficial to the fine-grained IE
 - Helps in Ontology learning

Current Trends in NE Research

- NER systems in non-newswire domains
 - Humanities (arts, history, archeology, literature etc.): *lots of non-traditional entities are present*
 - Chemical and bio-chemical (*long and nested NEs*)
 - Biomedical texts and clinical records (*no nomenclature and contains lots of common words*)
 - Unstructured datasets such as Twitter, online product reviews, blogs, SMS etc.

A brief introduction to Bio-text Mining

Aims: Text mining

- *Data Mining* -> needs structured data, usually in numerical form
- *Text mining*: discover & extract unstructured knowledge hidden in text—Hearst (1999)
- Text mining aids to construct hypotheses from associations derived from text
 - protein-protein interactions
 - associations of genes—*phenotypes*
 - functional relationships among genes...etc.

An Example

- *Stress is associated with migraines*
- *Stress can lead to loss of magnesium*

=> Loss of magnesium may cause migraine

Text Mining in biomedicine

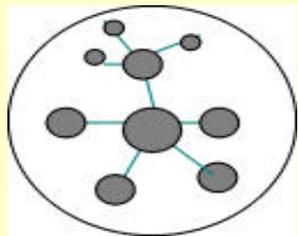
- Why biomedicine?
 - Consider just MEDLINE: 19,000,000 references, 40,000 added per month
 - Dynamic nature of the domain: new terms (*genes*, *proteins*, *chemical compounds*, *drugs* etc.) constantly created
 - Impossible to manage such an information overload

From Text to Knowledge:

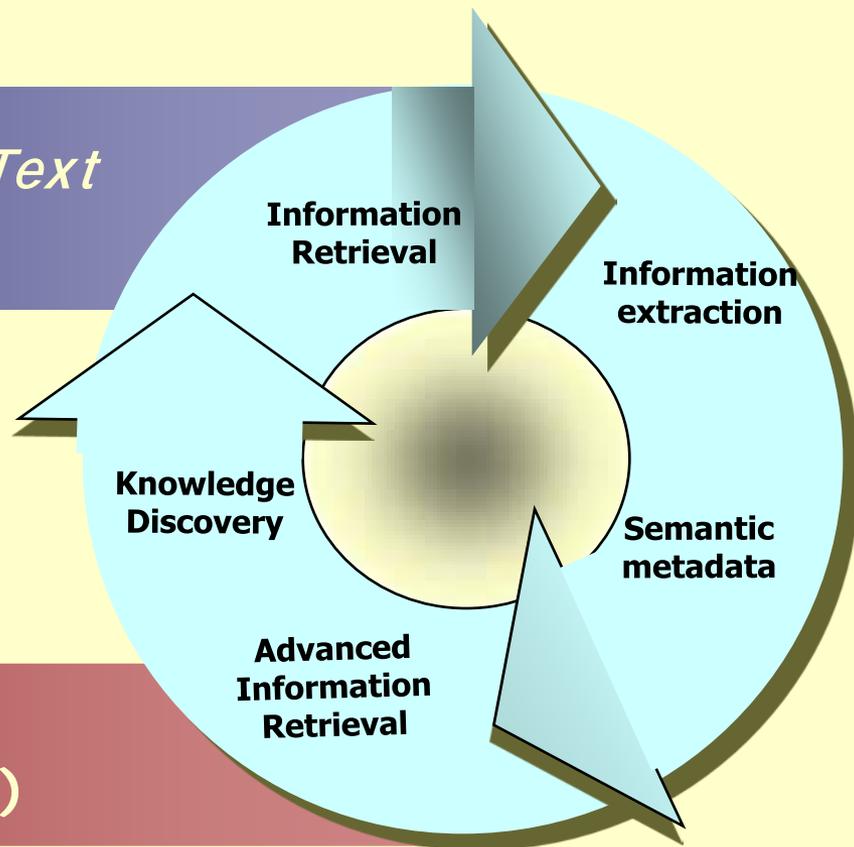
tackling the data deluge through text mining



Unstructured Text
(implicit knowledge)



Structured content
(explicit knowledge)



Some biotext mining campaigns

- KDD Cup-2002
- TREC-Genomics (<http://ir.ohsu.edu/genomics/>)
- JNLPBA-2004
(<http://www.nactem.ac.uk/tsujii/GENIA/ERtask/report.html>): Named entity recognition
- BioCreative (www.biocreative.org)- Information extraction including NER, PPI, text categorization etc. (2004, 2006, 2008 etc)
- BioNLP 2009, 2011, 2013-detailed biological phenomenon
(<http://www.nactem.ac.uk/tsujii/GENIA/SharedTask>)

Weighted vote based classifier
Ensemble (already discussed)

NE Extraction in Biomedicine

- **Objective**-identify biomedical entities and classify them into some predefined categories
 - *E.g. Protein, DNA, RNA, Cell_Line, Cell_Type*
- *Major Challenges*
 - building a complete dictionary for all types of biomedical NEs is infeasible due to the generative nature of NEs
 - NEs are made of very long compounded words (i.e., contain nested entities) or abbreviations and hence difficult to classify them properly
 - names do not follow any nomenclature

Challenges (Contd..)

- NEs include different symbols, common words and punctuation symbols, conjunctions, prepositions etc.
 - NE boundary identification is more difficult and challenging
- Same word or phrase can refer to different NEs based on their contexts

Features

- **Context Word:** Preceding and succeeding words
- **Word Suffix and Prefix**
 - **Fixed length** character strings stripped from the **ending** or **beginning** of word
- **Class label:** **Class label(s)** of the previous word (s)
- **Length (binary valued):** Check whether the length of the current word less than **three** or not (shorter words rarely NEs)
- **Infrequent (binary valued):** Infrequent words in the training corpus most probably NEs

Features

- **Part of Speech (PoS) information**- PoS of the current and/or surrounding token(s)
 - GENIA tagger V2.0.2 (<http://www-tsujii.is.s.u-tokyo.ac.jp/GENIA/tagger>)
- **Chunk information**-Chunk of the current and/or surrounding token(s)
 - GENIA tagger V2.0.2
- **Unknown token feature**-checks whether current token appears in training

Features

- Word normalization

- feature attempts to reduce a word to its stem or root form (from GENIA tagger O/P)

- Head nouns

- major noun or noun phrase of a NE that describes its function or the property
- E.g. *factor* is the head noun for the NE *NF-kappa B transcription factor*

Features

- **Verb trigger**-special type of verb (e.g., *binds*, *participates* etc.) that occur preceding to NEs and provide useful information about the NE class
- **Word class feature**-Certain kinds of NEs, which belong to the same class, are similar to each other
 - capital letters → A, small letters → a, number → O and non-English characters → -
 - consecutive same characters are squeezed into one character
 - groups similar names into the same NE class

Features

- Informative words
 - NEs are too long, complex and contain many common words that are actually not NEs
 - Function words- *of, and* etc.; nominals such as *active, normal* etc. appear in the training data often more frequently but these don't help to recognize NEs
 - Feature extracts informative words from training data statistically
- Content words in surrounding contexts-*Exploits global context information (feature extraction for the test set was based on the GENIA tagger)*

Features

- *Orthographic Features*-number of orthographic features depending upon the contents of the wordforms

| Feature | Example | Feature | Example |
|---------------|-----------------|-----------------|----------------|
| InitCap | Src | AllCaps | EBNA, LMP |
| InCap | mAb | CapMixAlpha | NFkappaB, EpoR |
| DigitOnly | 1, 123 | DigitSpecial | 12-3 |
| DigitAlpha | 2× NFkappaB, 2A | AlphaDigitAlpha | IL23R, EIA |
| Hyphen | - | CapLowAlpha | Src, Ras, Epo |
| CapsAndDigits | 32Dc13 | RomanNumeral | I, II |
| StopWord | at, in | ATGCSeq | CCGCCC, ATAGAT |
| AlphaDigit | p50, p65 | DigitCommaDigit | 1,28 |
| GreekLetter | alpha, beta | LowMixAlpha | mRNA, mAb |

Experiments

- Datasets-JNLPBA 2004 shared task datasets
 - Training: 2000 MEDLINE abstracts with 500K wordforms
 - Test: 404 abstracts with 200K wordforms
- Tagset: 5 classes
 - Protein, DNA, RNA, Cell_line, Cell_type
- Classifiers
 - CRF and SVM
- Evaluation scheme: JNLPBA 2004 shared task script (<http://www-tsujii.is.s.u-tokyo.ac.jp/GENIA/ERTask/report.html>)
 - Recall, precision and F-measure according to *exact boundary match*, *right* and *left* boundary matching

Experiments

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best individual classifier | 73.10 | 76.76 | 74.76 |
| Baseline-1 | 71.03 | 75.76 | 73.32 |
| Baseline-II | 71.42 | 75.90 | 73.59 |
| Baseline-III | 71.72 | 76.25 | 73.92 |
| SOO based ensemble | 74.17 | 77.87 | 75.97 |

- Baseline-I: Simple majority voting of the classifiers
- Baseline-II: Weighted voting where weights are based on the overall F-measure value
- Baseline-III: Weighted voting where weights are the F-measure of the individual classes

Issues of corpus compatibilities

Issues of Cross-corpus Compatibilities

- No unified annotation scheme exists for the biomedical entity annotation
- Building a system that performs reasonably well for almost all the domains is important!
- Datasets used in the experiments
 - JNLPBA shared task datasets
 - GENETAG datasets
 - AIMed datasets
- Differ in text selection as well as annotation

Experimental Setups

- **Experimental Setup-I:**
 - GENIA corpus by replacing all tags except 'Protein' by 'O' (other-than-NE) + AIMed corpus
 - Cross-validation

- **Experimental Setup-II:**
 - 'Protein' and 'DNA' annotations of GENIA+ Replace all other annotations by 'O' + AIMed corpus
 - Cross-validation

Experiments

- **Experimental Setup-III:**
 - GENIA corpus by replacing all tags except ‘Protein’ by ‘O’ (other-than-NE) + GENETAG corpus
 - Test on GENETAG
- **Experimental Setup-IV:**
 - GENIA with only ‘Protein’, ‘DNA’ and ‘RNA’ annotations + GENETAG corpus
 - Test on GENETAG corpus

Results: Cross Corpus

| Approach | Training set | Test set | Recall | Precision | F-measure |
|----------------------|----------------------------------|------------------|--------|-----------|-----------|
| Best Ind. Classifier | JNLPBA (protein only)+AIMed | AIMed | 83.14 | 83.19 | 83.17 |
| SOO | JNLPBA (protein only)+AIMed | AIMed | 85.10 | 85.01 | 85.05 |
| Best Ind. Classifier | JNLPBA (protein + DNA)+AIMed | AIMed | 82.17 | 84.15 | 83.15 |
| SOO | JNLPBA (protein + DNA)+AIMed | Cross validation | 84.07 | 86.01 | 85.03 |
| Best Ind. Classifier | JNLPBA (protein only)+GENETAG | GENETAG | 89.44 | 93.07 | 91.22 |
| SOO | JNLPBA (protein only)+GENETAG | GENETAG | 91.19 | 94.98 | 93.05 |
| Best Ind. Classifier | JNLPBA (protein+DNA+RNA)+GENETAG | GENETAG | 88.70 | 93.55 | 91.06 |
| SOO | JNLPBA (protein+DNA+RNA)+GENETAG | GENETAG | 90.09 | 95.16 | 92.56 |

Results: Original Datasets

| Dataset | Model | Recall | Precision | F-measure |
|---------|----------------------------|--------|-----------|-----------|
| GENIA | Best individual classifier | 73.10 | 76.78 | 74.90 |
| | SOO | 74.17 | 77.87 | 75.97 |
| AIMed | Best individual classifier | 94.56 | 92.66 | 93.60 |
| | SOO | 95.65 | 94.23 | 94.93 |
| GENETAG | Best individual classifier | 95.35 | 95.31 | 95.33 |
| | SOO | 95.99 | 95.81 | 95.90 |

Drop in performance by around 10% for AIMed
and around 3% for GENETAG

Ensemble Selection based on MOO

(already discussed)

Experiments

- Base classifiers
 - Based on different feature representations, several CRF and SVM classifiers built
- Objective functions (*in this work*)
 1. MOO1: overall average *recall* and *precision*
 2. MOO2: average F-measure value of five classes
 3. MOO3: average recall and precision values of five classes
 4. MOO4: average F-measure values of individual NE boundaries

Experiments (Results)

| Model | Recall | Precision | F-measure |
|----------------------------|--------|-----------|-----------|
| Best individual classifier | 73.10 | 76.78 | 74.90 |
| MOO1 | 75.52 | 78.03 | 76.75 |
| MOO2 | 75.78 | 78.45 | 77.09 |
| MOO3 | 75.91 | 78.98 | 77.41 |
| MOO4 | 76.15 | 79.09 | 77.59 |

Around 2% improvement over the present state-of-the-art

Experiments

| Class | | recall | precision | F-measure |
|-----------|------------------------|--------|-----------|-----------|
| Overall | FULLY correct | 76.78 | 73.10 | 74.90 |
| | correct LEFT boundary | 80.56 | 76.69 | 78.58 |
| | correct RIGHT boundary | 83.98 | 79.95 | 81.92 |
| Protein | FULLY correct | 82.31 | 73.22 | 77.50 |
| | correct LEFT boundary | 86.89 | 77.30 | 81.81 |
| | correct RIGHT boundary | 88.70 | 78.91 | 83.51 |
| cell_line | FULLY correct | 59.29 | 56.62 | 57.93 |
| | correct LEFT boundary | 64.31 | 61.41 | 62.82 |
| | correct RIGHT boundary | 71.68 | 68.45 | 70.03 |
| DNA | FULLY correct | 74.03 | 72.61 | 73.31 |
| | correct LEFT boundary | 76.46 | 75.00 | 75.72 |
| | correct RIGHT boundary | 81.17 | 79.62 | 80.39 |
| RNA | FULLY correct | 71.83 | 72.86 | 72.34 |
| | correct LEFT boundary | 74.65 | 75.71 | 75.18 |
| | correct RIGHT boundary | 80.28 | 81.43 | 80.85 |
| cell_type | FULLY correct | 69.21 | 78.95 | 73.76 |
| | correct LEFT boundary | 71.25 | 81.28 | 75.93 |
| | correct RIGHT boundary | 76.93 | 87.75 | 81.99 |

Reading

- Book on BioTextMining
 - S. Ananiadou & J. McNaught (eds) (2006) Text Mining for Biology and Biomedicine, ArtechHouse
 - McNaught, J. & Black, W. (2006) Information Extraction, Text Mining for Biology & Biomedicine, Artechhouse, pp.143-177
- Detailed bibliography in Bio-Text Mining
 - BLIMP <http://blimp.cs.queensu.ca/>
 - <http://www.ccs.neu.edu/home/futrelle/bionlp/>

***Thank you for your
attention!***

*Some of the slides have been taken
from...*

- Hamish Cunningham
- Beto Boullosa
- Sophia Ananiadou