Exploring Semantic Properties of Sentence Embeddings

Xunjie Zhu
Rutgers University

Tingfeng Li *
Northwestern Polytechnic University

Gerard de Melo
Rutgers University

*Work conducted while visiting Rutgers University

Introduction

- Sentence Embeddings:
  - Encode a variable-length input sentence into a constant size vector
- Examples:
  - Based on Word Embeddings:
    - (I) Glove Averaging (Wieting et al., 2015)
    - (II) Concatenated P-Mean Embeddings (R¨uckl´e et al. 2018)
    - (III) Sent2Vec (Pagliardini et al. 2018)
  - Based on RNNs:
    - (I) SkipThought Vectors (Kiros et al. 2015)
    - (II) InferSent (Conneau et al., 2017)

Goal

- Exploring what specific semantic properties are directly reflected by such embeddings.
- Focusing on a few select aspects of sentence semantics.
- Concurrent related work: Conneau et al. ACL 2018

  (i) Their work studies what you can learn to predict using 100,000 training instances

  (ii) Our goal: Directly study the embeddings (via cosine similarity)

Approach: Contrastive Sentences

Minor alterations of a sentence may lead to notable shifts in meaning.

(i) A rabbit is jumping over the fence (S)

(ii) A rabbit is hopping over the fence ($S^{=}$)

(iii) A rabbit is not jumping over the fence (S^{*})

Sentence Modification Schemes

- Not-Negation
- Quantifier-Negation
- Synonym Substitution
- Embedded Clause Extraction
- Passivization
- Argument Reordering
- Fixed Point Inversion

Negation Detection

- Original Sentence:
  - A person is slicing an onion.
- Synonym Substitution:
  - A person is cutting an onion.
- Not Negation:
  - A person is not slicing an onion.

Negation Variant

- Not Negation:
  - A man is not standing on his head under water.
- Quantifier Negation:
  - There is no man standing on his head under water.
- Original Sentence:
  - A man is standing on his head under water.

Clause Relatedness

- Original Sentence:
  - Octel said the purchase was expected.
- Extracted Clause:
  - The purchase was expected.
- Not Negation:
  - Octel said the purchase was not expected

Argument Sensitivity

- Original Sentence:
  - Francesca teaches Adam to adjust the microphone on his stage
- Passivization:
  - Adam is taught to adjust the microphone on his stage
- Argument Reordering:
  - Adam teaches Francesca to adjust the microphone on his stage

Fixed Point Reordering

- Original Sentence:
  - A black dog in the snow is jumping off the ground and catching a stick.
- Synonym Substitution:
  - A black dog in the snow is leaping off the ground and catching a stick.
- Fixed Point Inversion(Corrupted Sentence):
  - In the snow is jumping off the ground and catching a stick a black dog.

Models and Dataset

	Dataset	Embedding Dim
Glove Avg	Common Crawl	300
P Means	Common Crawl	300
Sent2Vec	English Wiki	600
SkipThought	Book Corpus	600
InferSent	SNLI	4096

	# of Sentences	From
Negation Detection	674	SICK, SNLI
Negation Variant	516	SICK, SNLI
Clause Relatedness	567	Penn Treebank MSR Paraphrase
Argument Sensitivity	445	SICK, MS Paraphrase
Fixed Point Reordering	623	SICK

Negation Detection

- Average of Word Embeddings is more easier misled by negation.

- Both InferSent and SkipThought succeed in distinguishing unnegated sentences from negated ones.

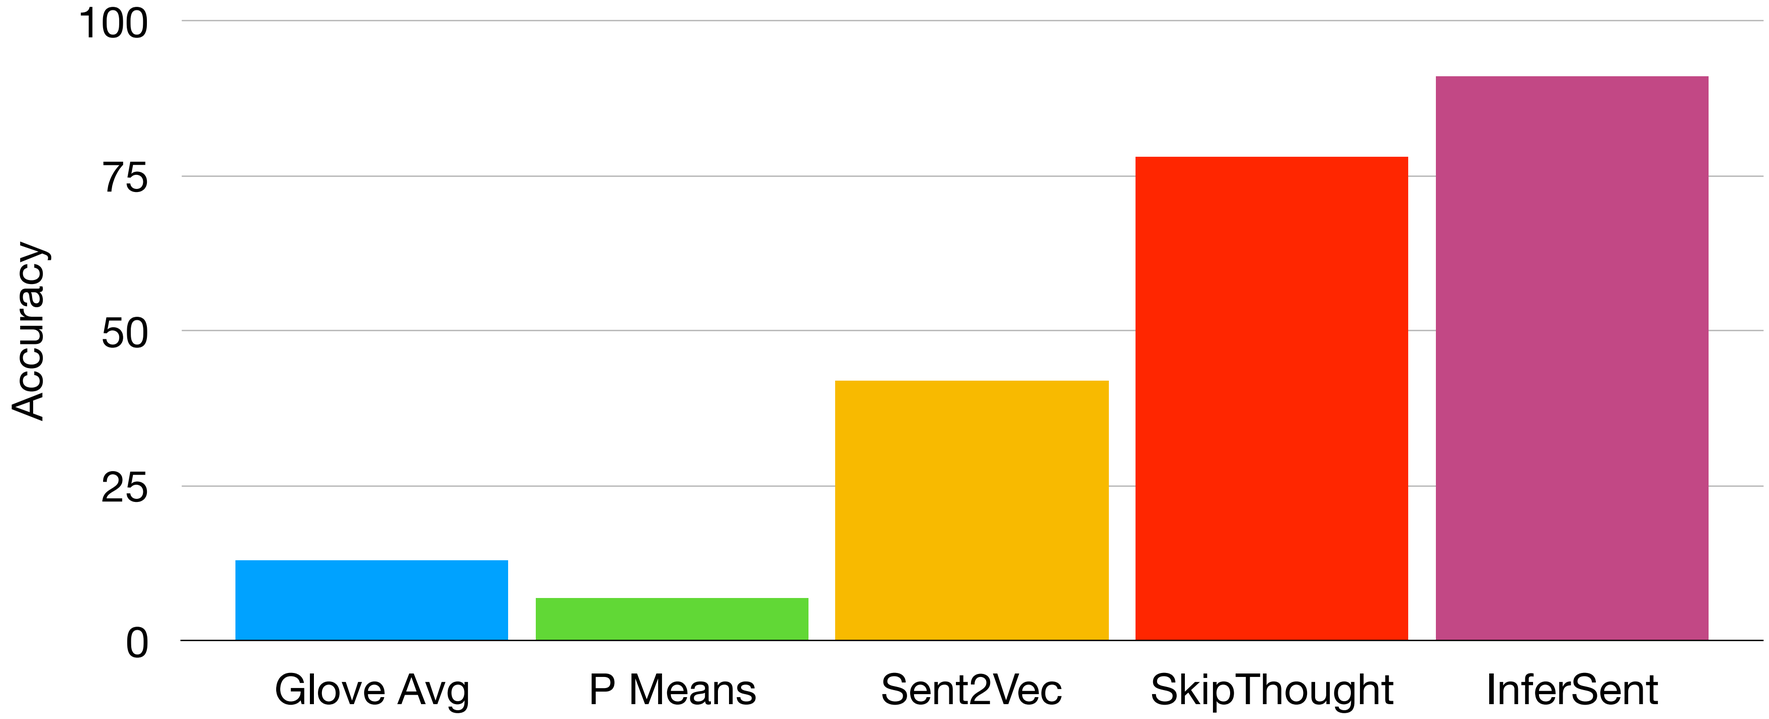

Zhu, Li & de Melo. Exploring Semantic Properties of Sentence Embeddings

RUTGERS
UNIVERSITY | NEW BRUNSWICK

Negation Variant

- Both averaging of word embeddings and SkipThought are dismal in terms of the accuracy.
- InferSent appears to have acquired a better understanding of negation quantifiers, as these are commonplace in many NLI datasets.

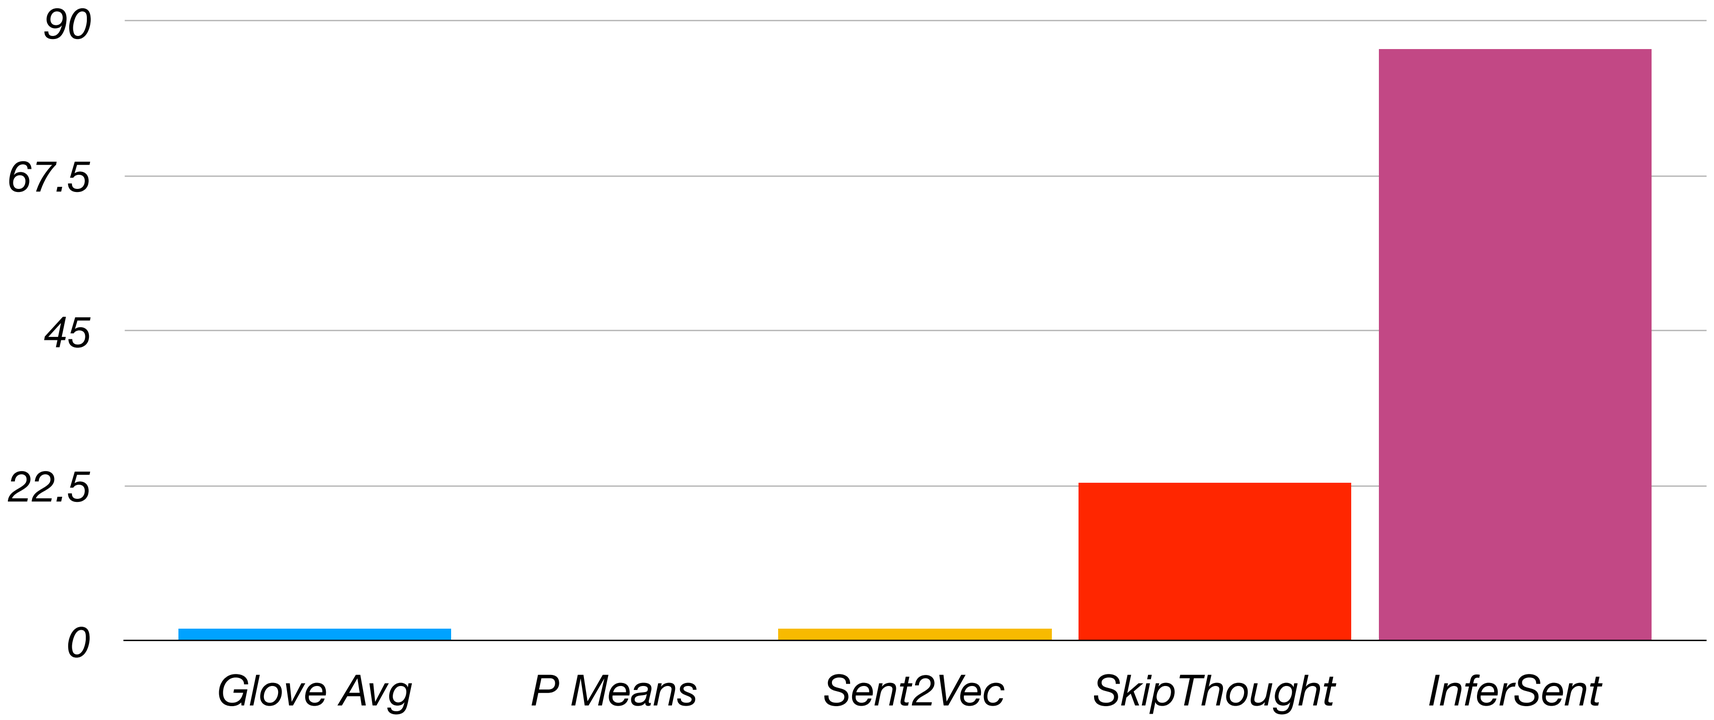

Zhu, Li & de Melo. Exploring Semantic Properties of Sentence Embeddings

Clause Relatedness

- Both SkipThought vectors and InferSent works poorly when sub clause is much shorter than original one.

- Sent2vec best in distinguishing the embedded clause of a sentence from a negation of that sentence.

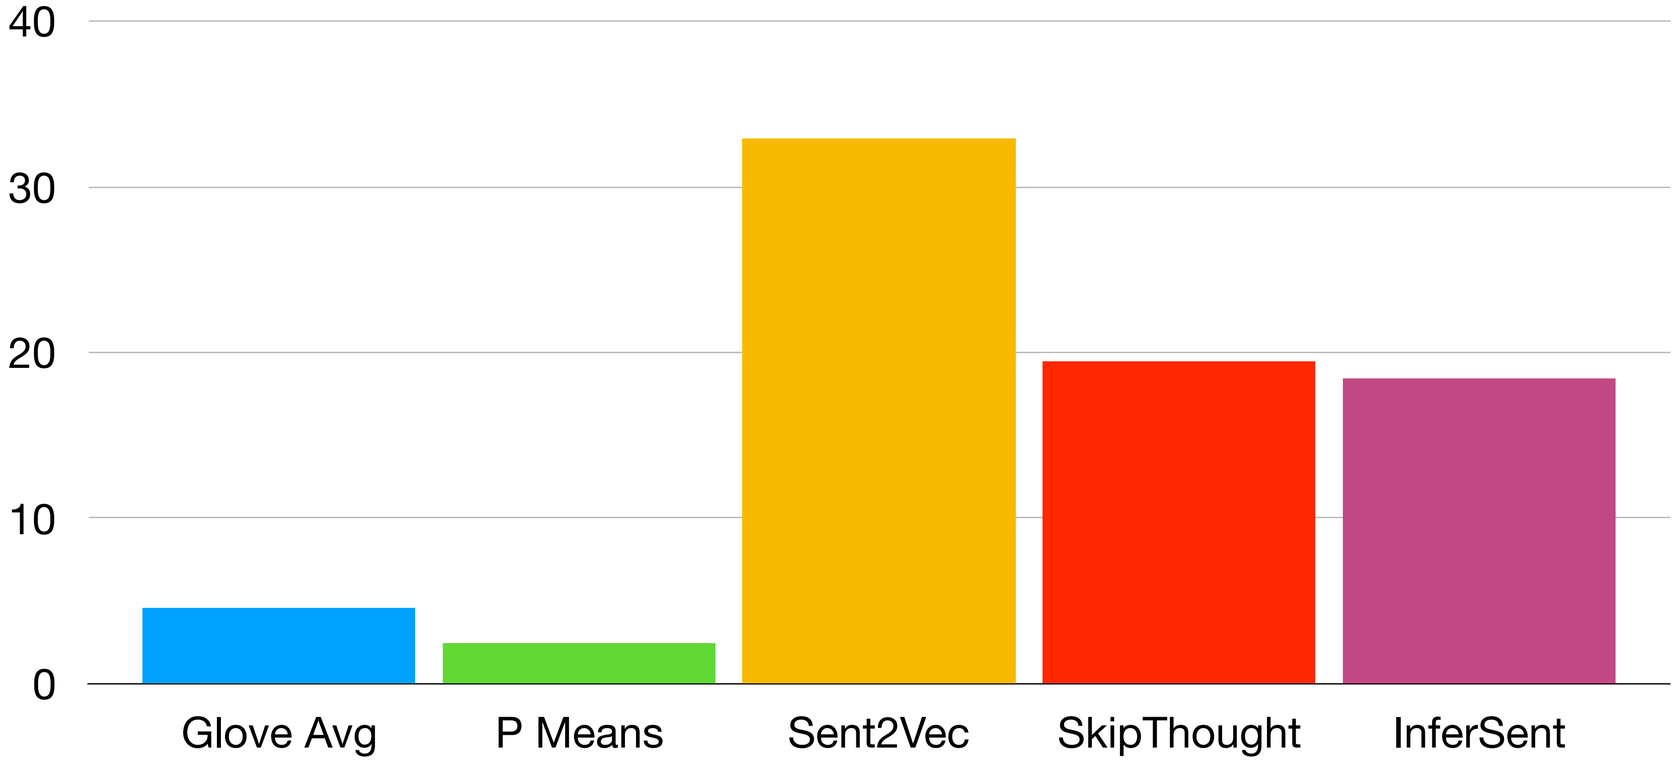

Zhu, Li & de Melo. Exploring Semantic Properties of Sentence Embeddings

RUTGERS
UNIVERSITY | NEW BRUNSWICK

Argument Sensitivity

- None of the analyzed approaches prove adept at distinguishing the semantic information from structural information in this case.

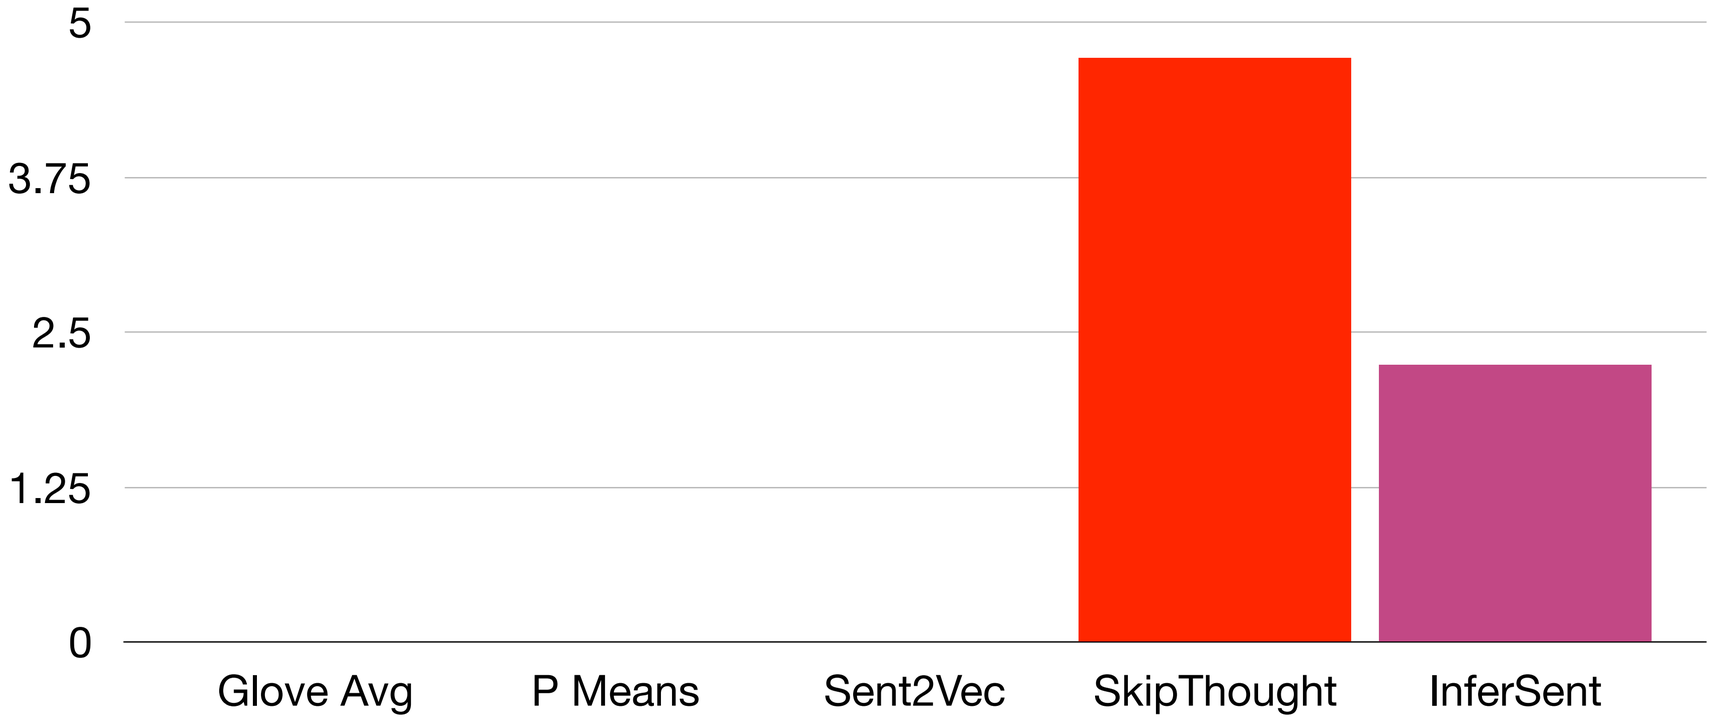

Zhu, Li & de Melo. Exploring Semantic Properties of Sentence Embeddings

RUTGERS
UNIVERSITY | NEW BRUNSWICK

Fixed Point Reordering

- Methods based on word embeddings do not encode sufficient word order information into the sentence embeddings.

- SkipThought and InferSent did well when the original sentence and its semantically equivalence share similar structure

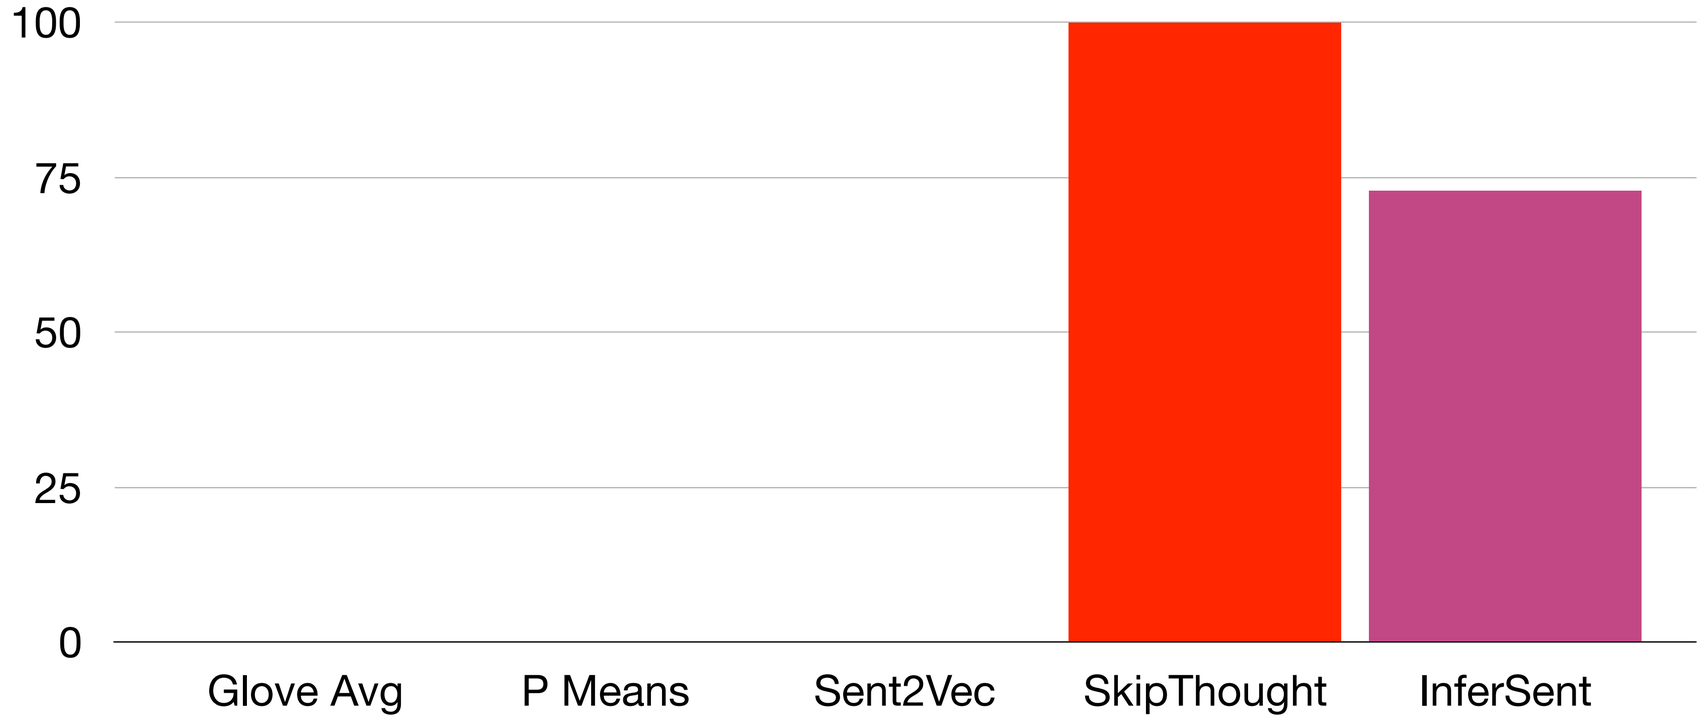

Zhu, Li & de Melo. Exploring Semantic Properties of Sentence Embeddings

Conclusion

- RNN based sentence embeddings better at identifying negation compared with word embedding based models
- Both SkipThought and InferSent distinguish negation of a sentence from synonymy.
- InferSent better at identifying semantic equivalence regardless of the order of words and copes better with quantifiers.
- SkipThoughts is more suitable for tasks in which the semantics of the sentence corresponds to its structure

Thank you!

Questions?
Contact us at
xunjie.zhu@rutgers.edu
and gdm@demelo.org