Character-based Kernels for Novelistic Plot Structure

Micha Elsner

School of Informatics
University of Edinburgh

April 27, 2012

We have good models for short articles...
Less research on ***storytelling***.

A ***use of language*** we don't understand...
Lots of ***data*** our tools don't cover.

Challenges

Storytelling is **typical** language use...

- But still little formal understanding of what a story ***is***!

Potential applications:

- Generating stories (games, education, etc.)
- Summarizing, searching, recommendations
- Sociolinguistics (fanfiction, NaNoWriMo)
  - $>$500k Harry Potter fan stories on fanfiction.net alone!

When good summarizers go bad...

...follows the main character Elizabeth Bennet as she deals with issues of manners, upbringing, morality, education and marriage... (Wikipedia)

The story turns on the marriage prospects of the five daughters of Mr. and Mrs. Bennet... (Amazon.com)

When good summarizers go bad...

...follows the main character Elizabeth Bennet as she deals with issues of manners, upbringing, morality, education and marriage... (Wikipedia)

The story turns on the marriage prospects of the five daughters of Mr. and Mrs. Bennet... (Amazon.com)

“Bingley.” Elizabeth felt Jane’s pleasure. “Miss Elizabeth Bennet.” Elizabeth looked surprised. “FITZWILLIAM DARCY” Elizabeth was delighted. Elizabeth read on: Elizabeth smiled. “If! “Dearest Jane! (Jason Huff: Microsoft Word ‘08)

Two approaches

Start small...

Sophisticated representations, simple texts:

- Generation: (McIntyre+Lapata) and others
- Annotation: (Volkova+al), Scheherezade (Elson+al)
- Analysis: AESOP (Goyal+al)

Start simple...

Complex texts, simple systems:

- Social networks: (Elson+Dames+McKeown)
- No-spoiler summaries: (Kazantseva+Szpakowicz)
- This project

Building something simple

- Basic enough to be robust...
- But not trivial...
- Experimental framework for comparisons

Similarity between novels

- Helpful for information retrieval:
  - Find another novel like “Pride and Prejudice”.
- Clustering and organization:
  - Are there “plot type” clusters?
- Project knowledge about training novels to unknown:
  - This novel is *like* “Pride and Prejudice”; maybe it’s a romance.

Overview

Motivation

Our Representation

Character, emotion and time

Implementation details

Preprocessing and graph kernels

Experiments

Telling real novels from surrogates

Analysis and conclusions

Plot is *high-level*...

Two basic insights:

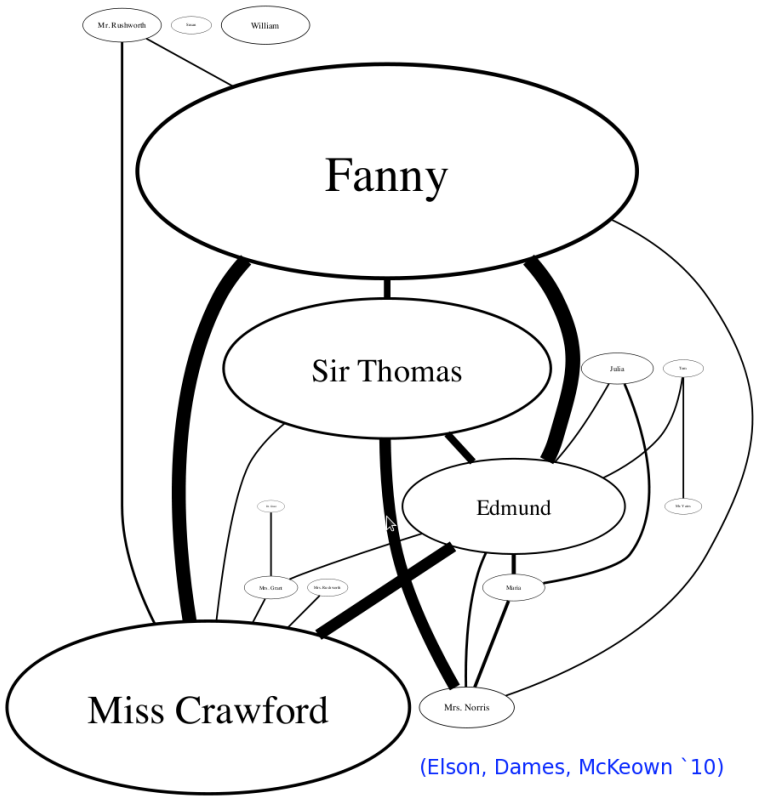

(Elson, Dames, McKeown `10)

Characters...
forming a *social network*

(Elson+al '10)

Plot is *high-level*...

Two basic insights:

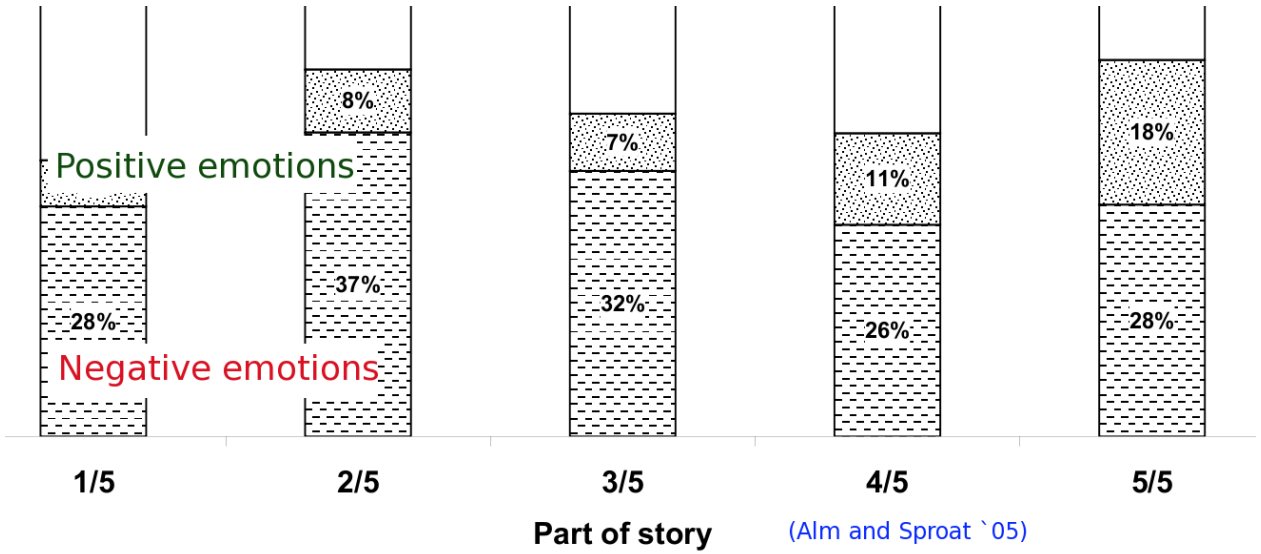

Story has an *emotional trajectory*

(Alm+Sproat '05)

Combine the two:

- Compute a trajectory for each character
- Observe social relationships through time

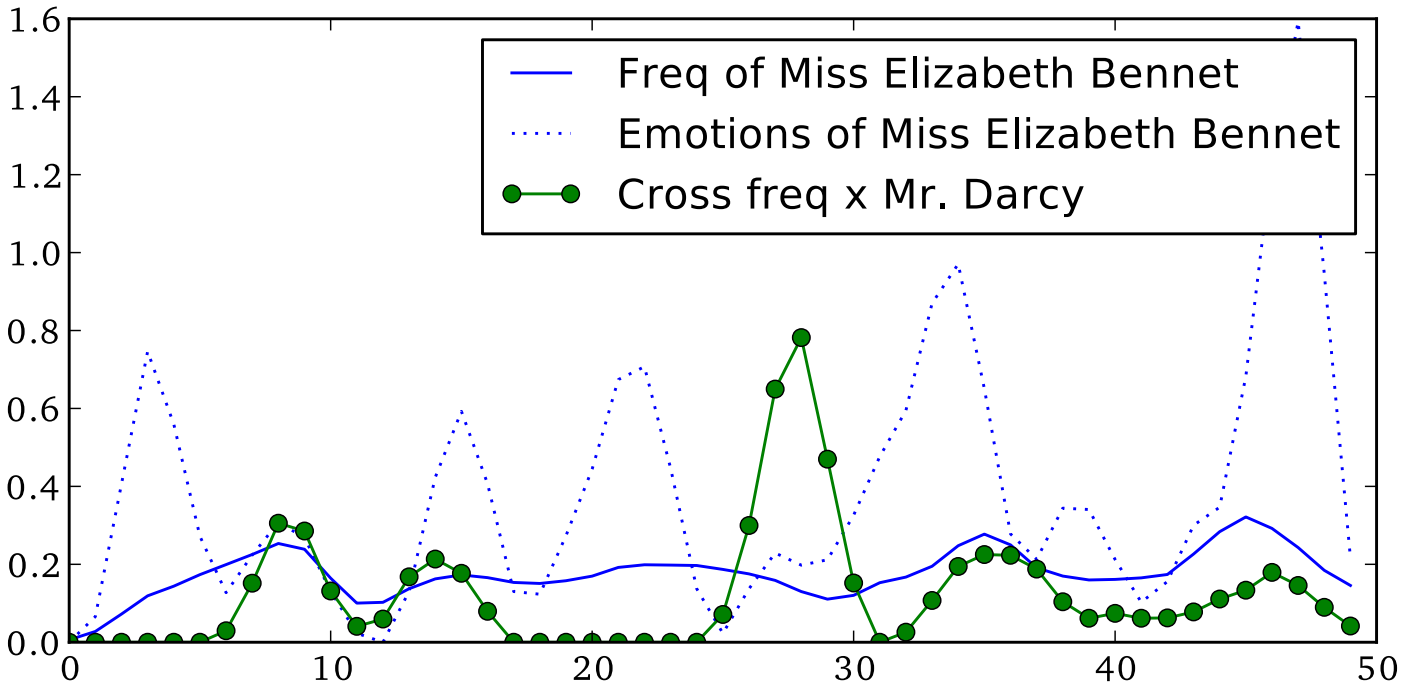

Combine the two:

- Compute a trajectory for each character
- Observe social relationships through time

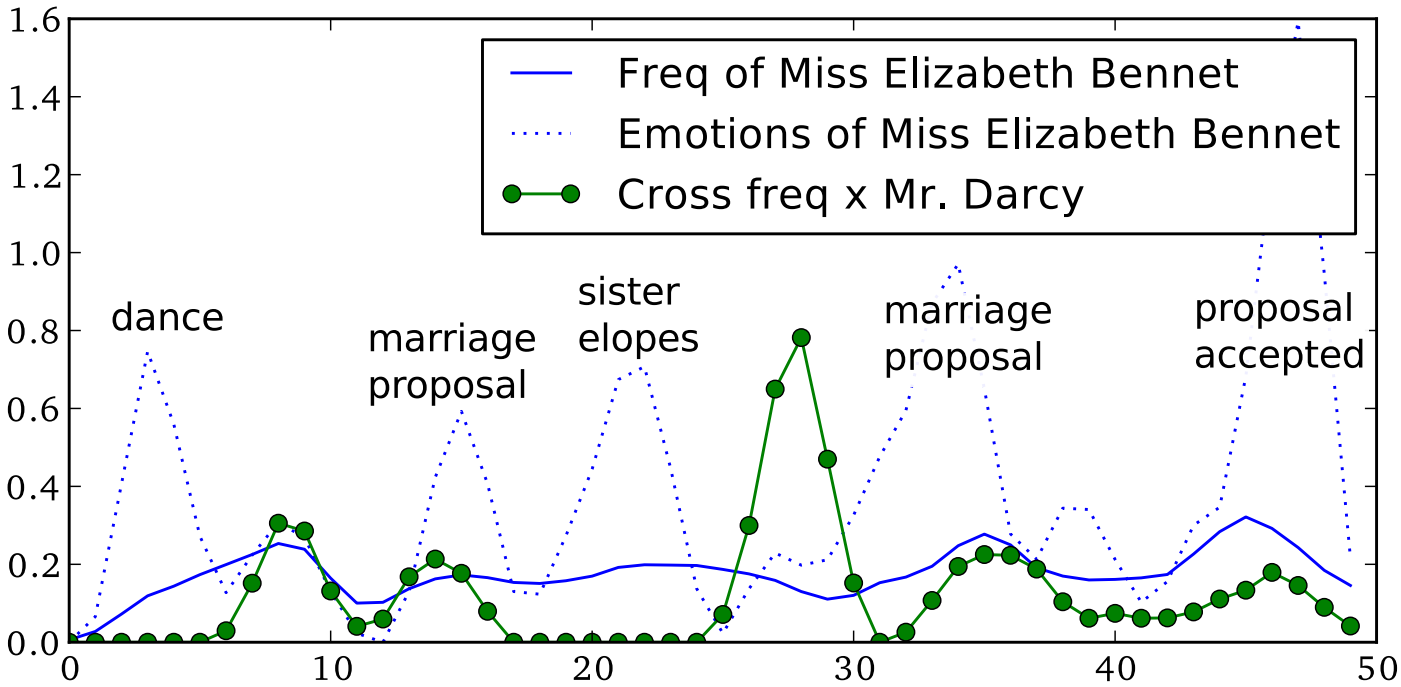

Preprocessing

- Chop the novel into paragraphs
- Parse everything and retrieve proper NPs
- Simple coreference on the NPs to find characters
- Emotion: “strong sentiment” cues from (Wilson+al ‘05)

Coreference

Similar to cross-document coreference:

- Shared name elements
- Presence in same documents
- List of gendered names and titles

“Miss Elizabeth Bennet” (f)	Elizabeth Bennet Elizabeth Miss Elizabeth Bennet Miss Bennet
“Miss Eliza” (f)	Miss Eliza Eliza
“Miss Elizabeth” (f)	Miss Elizabeth
“Lizzy” (?)	Lizzy

Use this representation to measure similarity...

Kernel function

$k(x, y)$: similarity between x and y

0: no similarity; > 0: more similar

basic ML building block

Use this representation to measure similarity...

Kernel function

$k(x, y)$: similarity between x and y
0: no similarity; > 0: more similar
basic ML building block

Use *convolution theorem* (Haussler '99) to build a complex kernel out of simpler ones:

$$k(x, y) = \sum_{ch_1 \in X} \sum_{ch_2 \in Y} \underbrace{c(ch_1, ch_2)}_{\text{kernel over characters}}$$

Similarity between characters

$e(ch_1, ch_2)$:

- Similarity for trajectory curves
- Normalized integral of the product
- Used for frequency and emotion

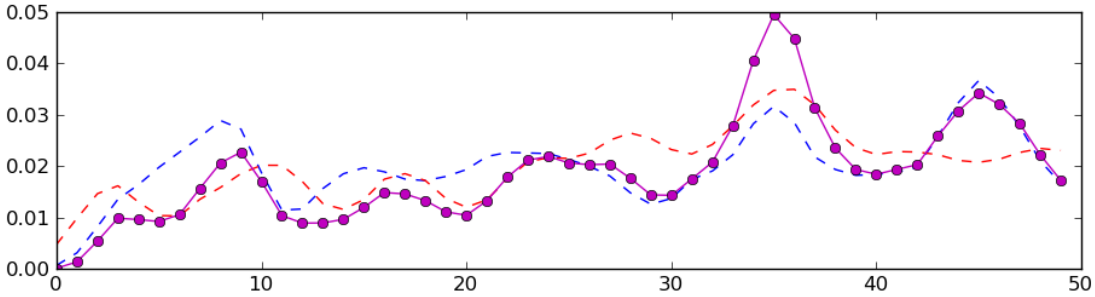

$d(ch_1, ch_2)$

- Nearby words

replied Elizabeth	17
Elizabeth felt	14
Elizabeth looked	10
Elizabeth's mind	7
...	

First-order character kernel

$$c_1(ch_1, ch_2) = d(ch_1, ch_2)e(ch_1, ch_2)$$

Adding social network features

Characters are more similar if:

- They each have close friends...
  - (Measured by co-occurrence frequency)
- ...who are also similar

Second-order character kernel

$$c_2(ch_1, ch_2) = c_1(ch_1, ch_2) \sum_{u' \in x} \sum_{v' \in Y} \underbrace{e(\widehat{u, u'}, \widehat{v, v'})}_{\text{relationship strength}} c_1(u', v')$$

Testing similarity

- First, simple proof of concept
- Independent of particular critical theory
- Difficult for very naive models

Testing similarity

- First, simple proof of concept
- Independent of particular critical theory
- Difficult for very naive models

Order discrimination

(Karamanis+al '04) (Barzilay+Lapata '05)

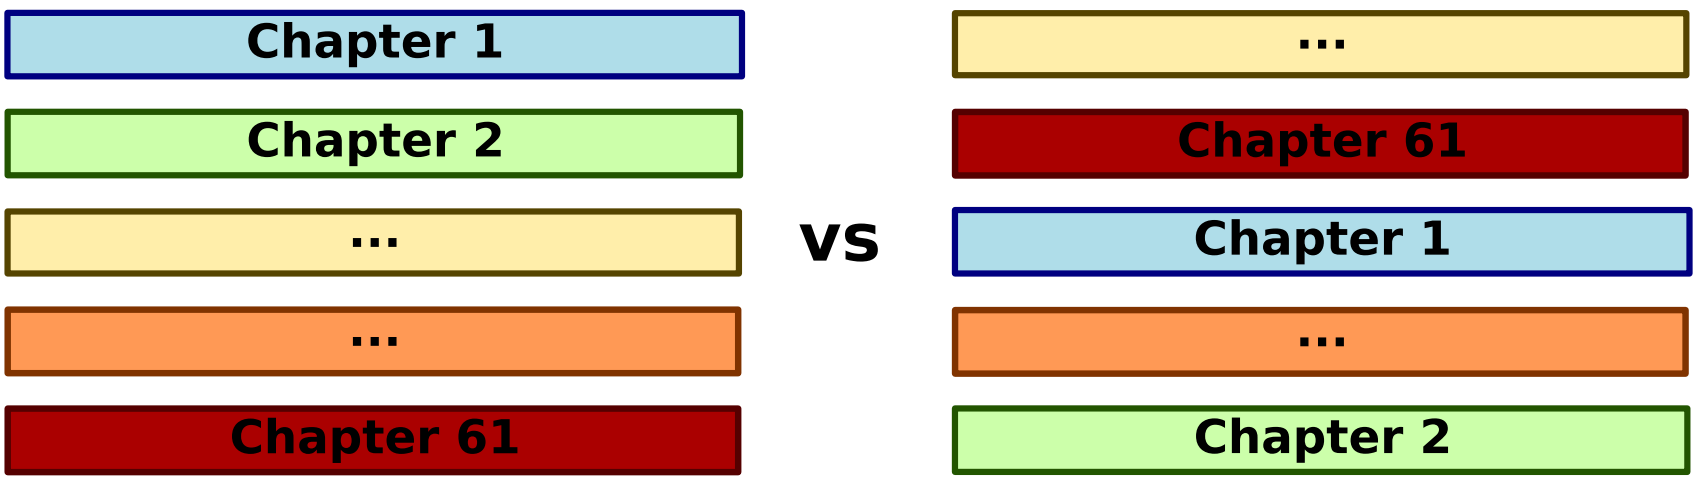

Weighted nearest-neighbor

For training set T, is:

$$\sum_{t \in T} k(t, y) > \sum_{t \in T} k(t, y_{perm})?$$

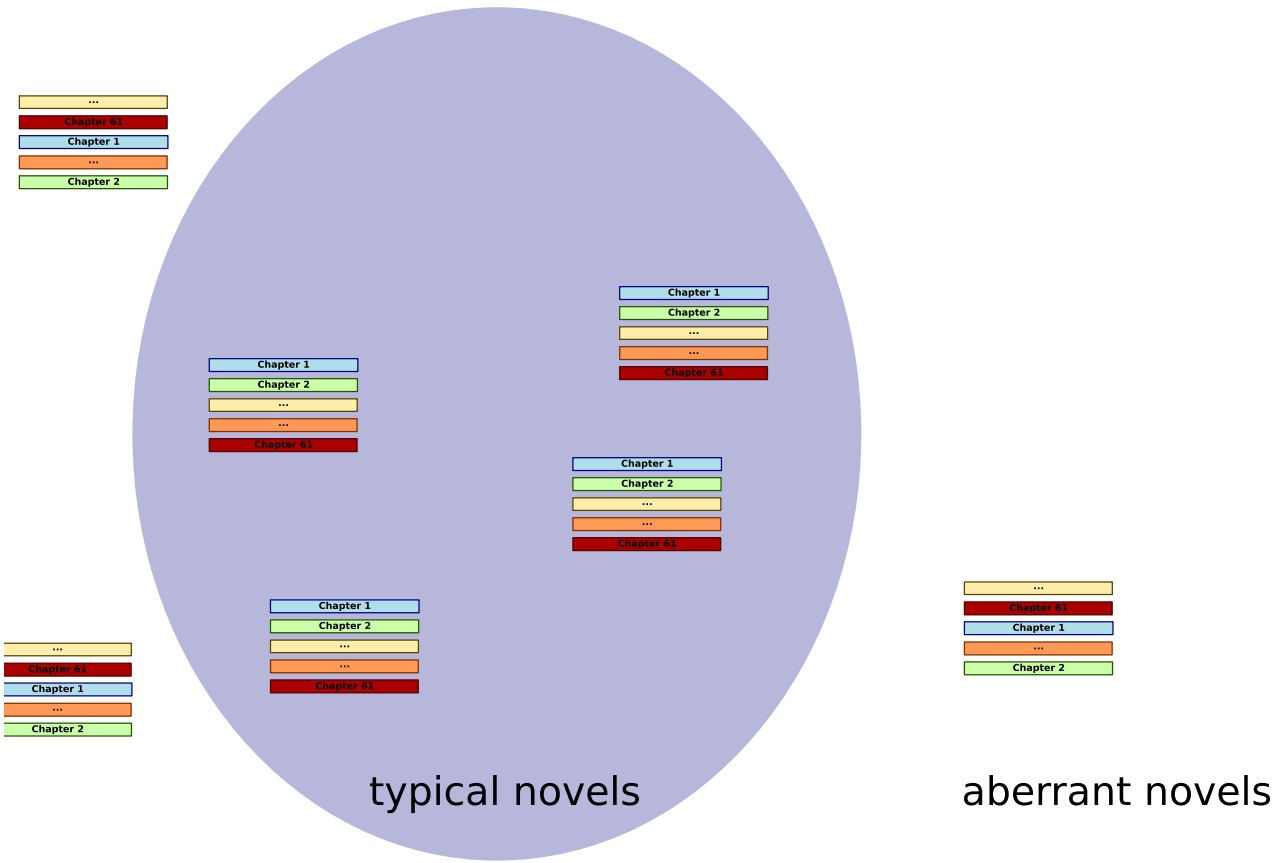

Weighted nearest-neighbor

For training set T, is:

$$\sum_{t \in T} k(t, y) > \sum_{t \in T} k(t, y_{perm})?$$

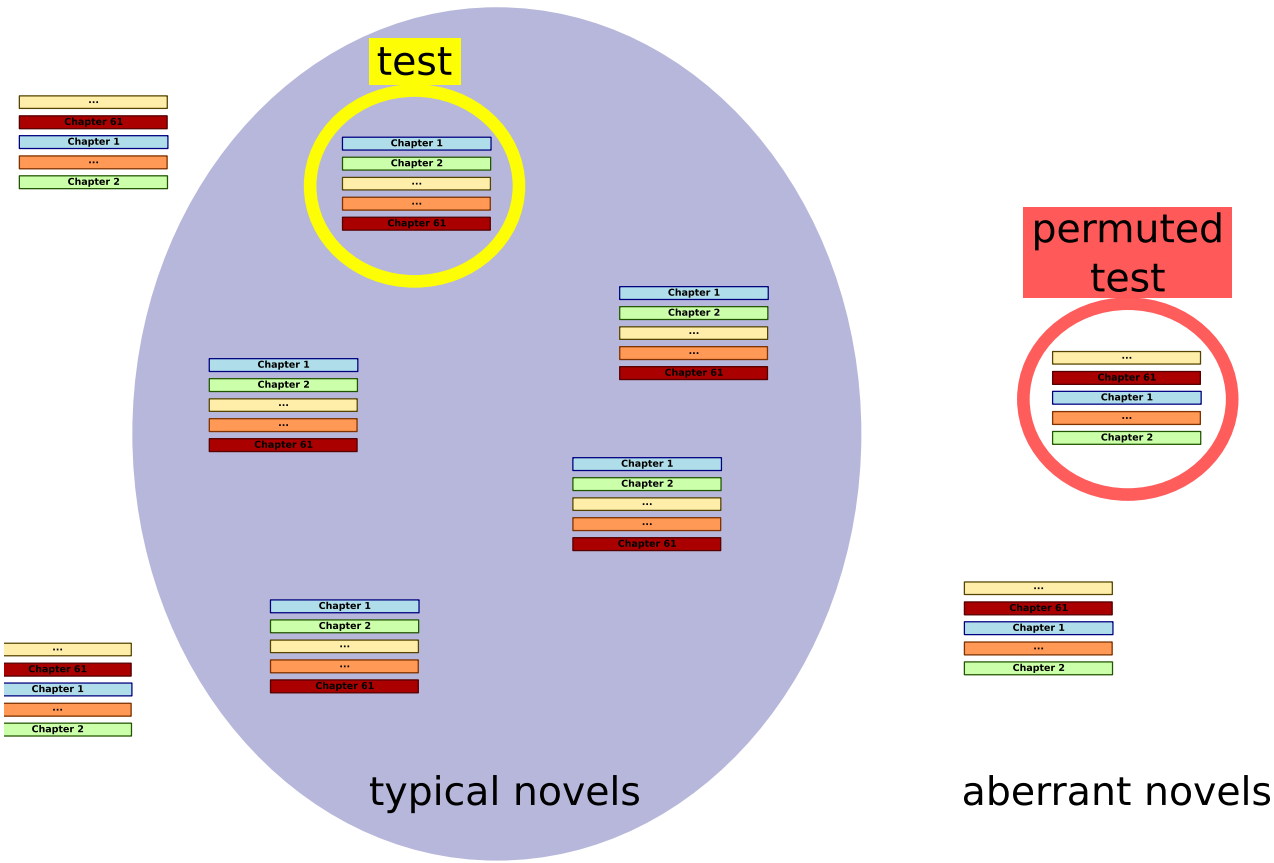

Weighted nearest-neighbor

For training set T, is:

$$\sum_{t \in T} k(t, y) > \sum_{t \in T} k(t, y_{perm})?$$

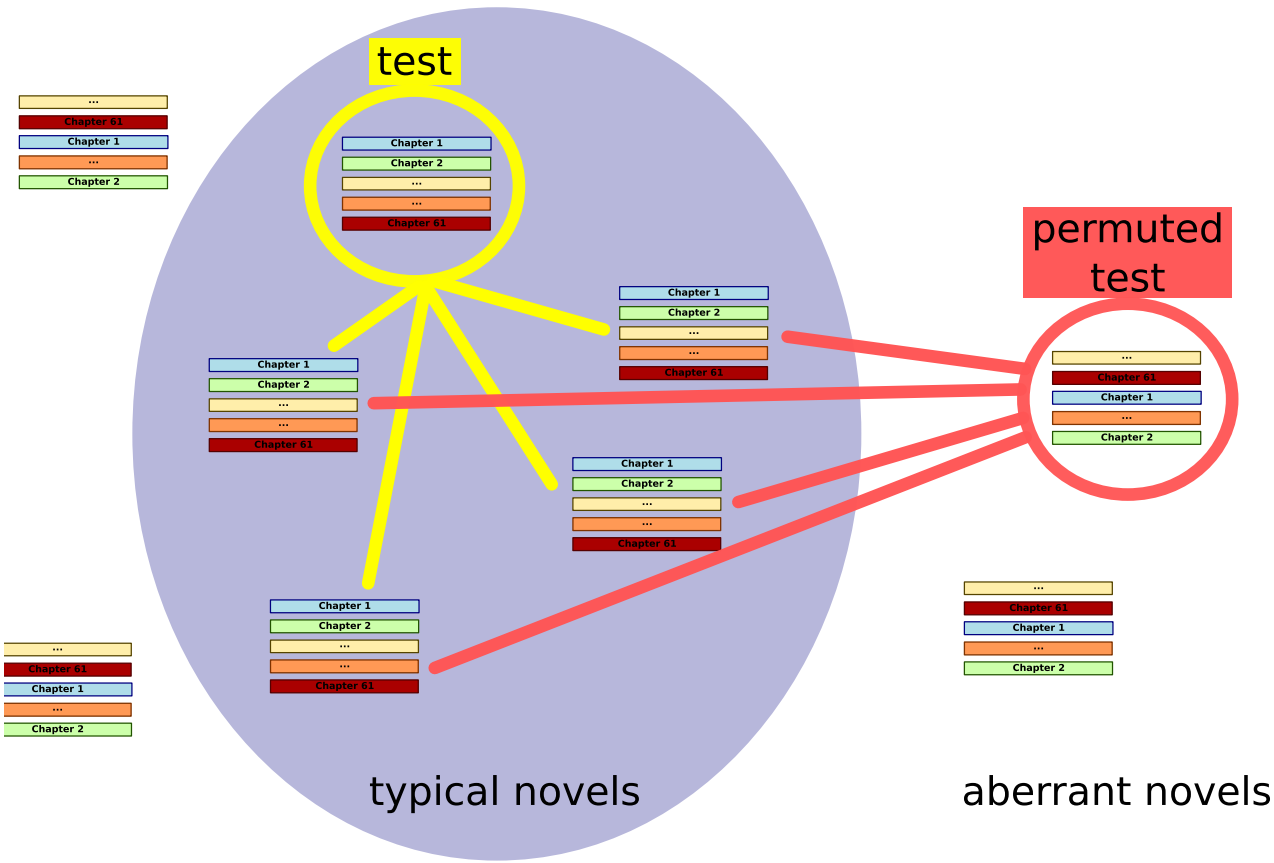

Results

(30 19th.c novels from Project Gutenberg)

Binary classifications

Chance accuracy 50%
Significance via kernel-based non-parametric test (Gretton+al '07)

	Random perm	Reversed

Results

(30 19th.c novels from Project Gutenberg)

Binary classifications

Chance accuracy 50%
Significance via kernel-based non-parametric test (Gretton+al '07)

	Random perm	Reversed
Whole-novel traj.	50	53

Results

(30 19th.c novels from Project Gutenberg)

Binary classifications

Chance accuracy 50%
Significance via kernel-based non-parametric test (Gretton+al '07)

	Random perm	Reversed
Whole-novel traj.	50	53
First-order k_1	77	63

Results

(30 19th.c novels from Project Gutenberg)

Binary classifications

Chance accuracy 50%
Significance via kernel-based non-parametric test (Gretton+al '07)

	Random perm	Reversed
Whole-novel traj.	50	53
First-order k_1	77	63
Second-order k_2	**90**	**67**

Character similarity

Hand-labeled characters from four Austen novels:

- Male, female, plural (“the Crawfords”), not a character (“London”)
- Protagonist, marriageable, other
- Ad-hoc scheme, probably doesn’t generalize much

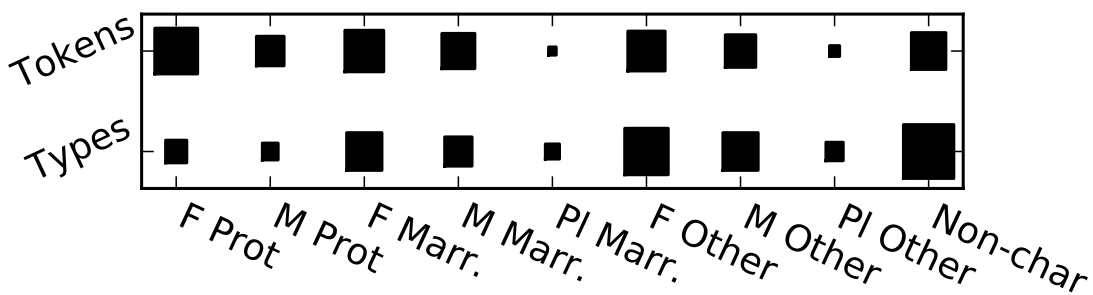

Character frequency by category

Similarity, unigram features only

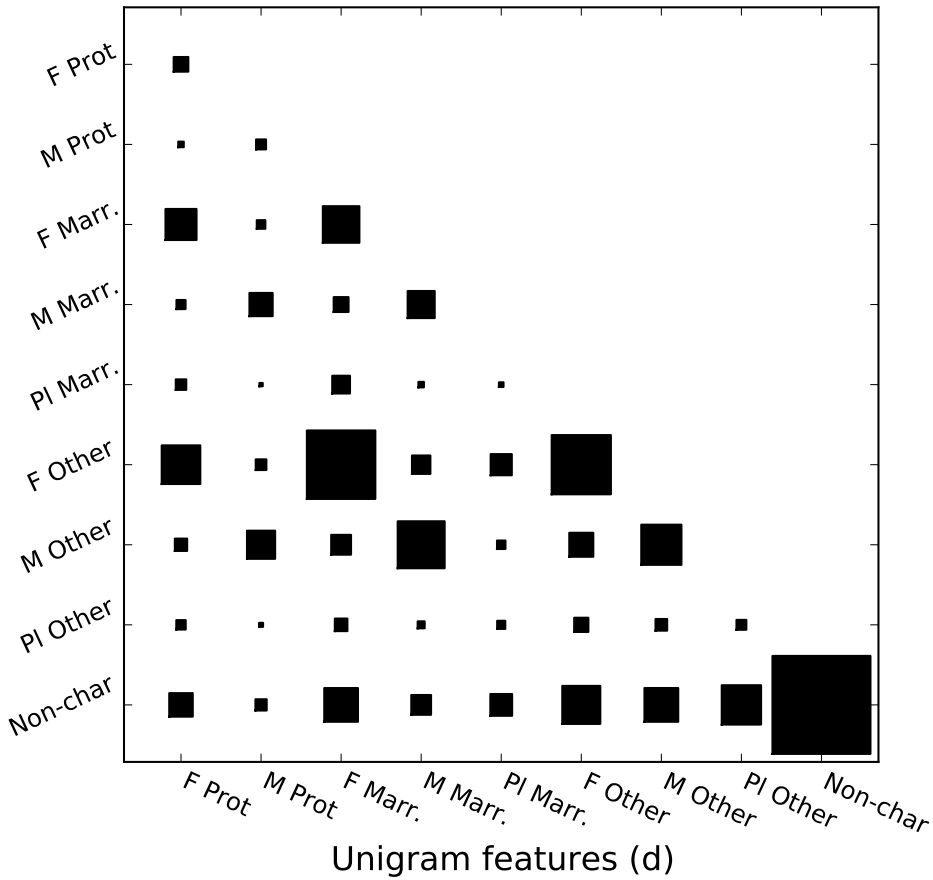

Similarity dependent mostly on non-characters!

Similarity, first-order kernel

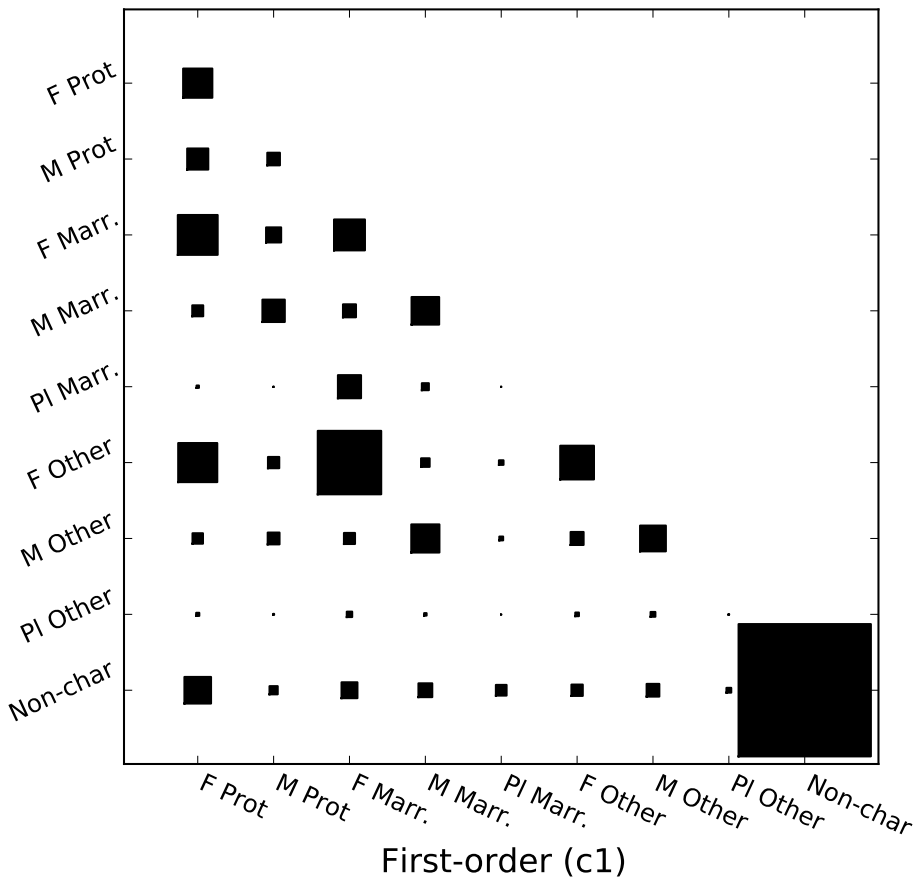

Non-characters still matter...
But similarity to real characters reduced.

Similarity, second-order kernel

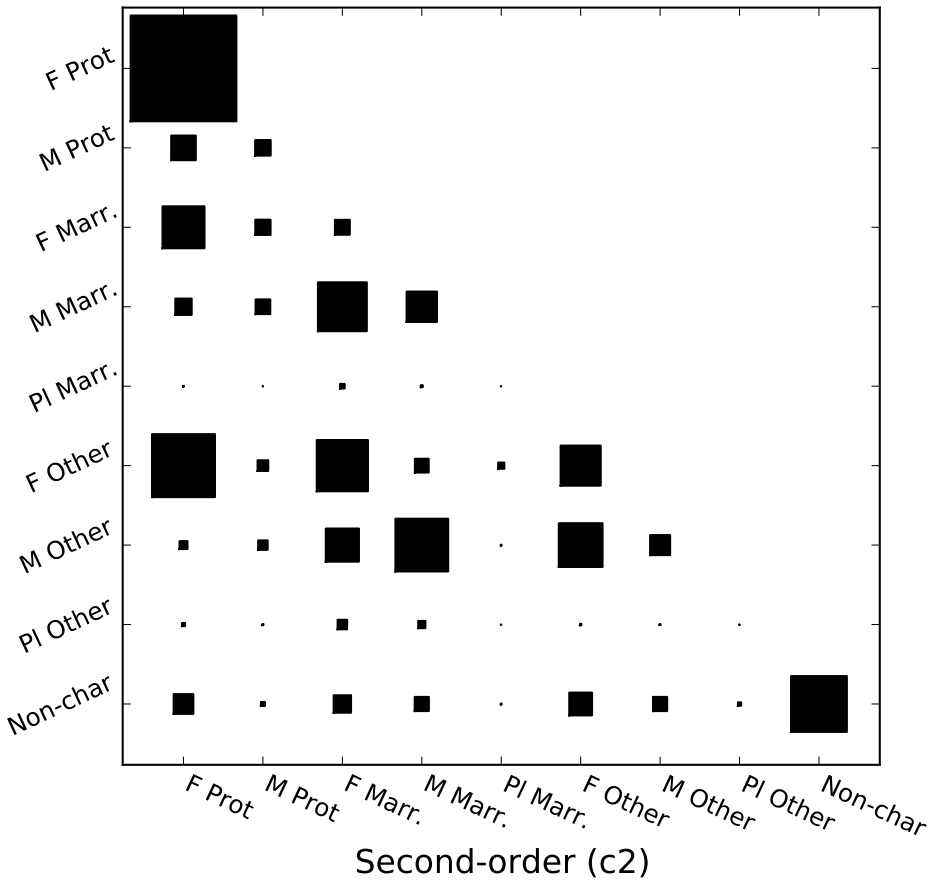

Characters with relationships more significant

Conclusions

- Plot structure: based on *character* and *emotion* over *time*
- Simple ordering test as proof of concept

Future work

- Eventually: search and summarize stories
- Topic modeling: match emotions to lexical features
- Interface for writers to visualize their work

Thanks: Sharon Goldwater, Mirella Lapata, Victoria Adams, Kira Mourão, and all of you!