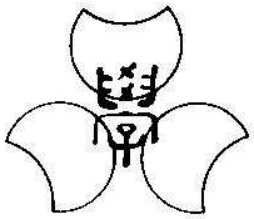

# Evidence for Gradient Salience: What Happens with Competing Non-salient Referents during Pronoun Resolution?

Ralph L. Rose
<rose@gpwu.ac.jp>
Gunma Prefectural Women's University
Gunma, Japan

30 November 2006
Australasian Language Technology Workshop
University of Sydney, Australia

# Introduction

Many practical implementations of pronoun resolution models rank candidate antecedents in a gradient manner (cf., Lappin and Leass, 1994; Tetreault, 2001).

# Introduction

Many practical implementations of pronoun resolution models rank candidate antecedents in a gradient manner (cf., Lappin and Leass, 1994; Tetreault, 2001).

Many psycholinguistic models similarly assume gradient ranking of antecedents (cf., Gernsbacher and Hargreaves, 1988; Hudson-D'Zmura and Tanenhaus, 1997). After many researchers, I'll refer to this ranking as *salience*.

# Introduction

Many practical implementations of pronoun resolution models rank candidate antecedents in a gradient manner (cf., Lappin and Leass, 1994; Tetreault, 2001).

Many psycholinguistic models similarly assume gradient ranking of antecedents (cf., Gernsbacher and Hargreaves, 1988; Hudson-D'Zmura and Tanenhaus, 1997). After many researchers, I'll refer to this ranking as *salience*.

Gradient salience has not actually been tested. Results could be explained in terms of a categorical model of salience.

# Introduction

Many practical implementations of pronoun resolution models rank candidate antecedents in a gradient manner (cf., Lappin and Leass, 1994; Tetreault, 2001).

Many psycholinguistic models similarly assume gradient ranking of antecedents (cf., Gernsbacher and Hargreaves, 1988; Hudson-D'Zmura and Tanenhaus, 1997). After many researchers, I'll refer to this ranking as *salience*.

Gradient salience has not actually been tested. Results could be explained in terms of a categorical model of salience.

| Do people actually use a gradient approach to salience ranking in pronoun reference resolution? |
|---|

# Overview

- Pronoun Reference Resolution
- Computing Salience of Referents
- Experiment
- Results and Analysis
- Discussion
- Further Work

# Overview

- Pronoun Reference Resolution
- Computing Salience of Referents
- Experiment
- Results and Analysis
- Discussion
- Further Work

# Overview

- Pronoun Reference Resolution
- Computing Salience of Referents
- Experiment
- Results and Analysis
- Discussion
- Further Work

# Overview

- Pronoun Reference Resolution
- Computing Salience of Referents
- Experiment
- Results and Analysis
- Discussion
- Further Work

# Overview

- Pronoun Reference Resolution
- Computing Salience of Referents
- Experiment
- Results and Analysis
- Discussion
- Further Work

# Overview

- Pronoun Reference Resolution
- Computing Salience of Referents
- Experiment
- Results and Analysis
- Discussion
- Further Work

# Pronoun Reference Resolution

When a pronoun is encountered in text, the processing faculty seeks to interpret the pronoun as referring to an entity introduced in previous discourse. These entities are often called *discourse referents* (Karttunen, 1976; Heim, 1982).

# Pronoun Reference Resolution

When a pronoun is encountered in text, the processing faculty seeks to interpret the pronoun as referring to an entity introduced in previous discourse. These entities are often called *discourse referents* (Karttunen, 1976; Heim, 1982).

The processor searches these referents and selects the most salient referent as the default interpretation of the pronoun.

# Pronoun Reference Resolution

When a pronoun is encountered in text, the processing faculty seeks to interpret the pronoun as referring to an entity introduced in previous discourse. These entities are often called *discourse referents* (Karttunen, 1976; Heim, 1982).

The processor searches these referents and selects the most salient referent as the default interpretation of the pronoun.

This basic description is prevalent in the literature:

- Psycholinguistic Models: Almor (1999); Greene et al. (1992)
- Computational Implementations: Hirst (1981); Mitkov (2002)

# Pronoun Reference Resolution

Consider two approaches to salience ranking:

- categorical salience: referents are salient or non-salient
- gradient salience: referents are salient to varying degrees

# Pronoun Reference Resolution

Consider two approaches to salience ranking:

- categorical salience: referents are salient or non-salient
- gradient salience: referents are salient to varying degrees

Most psycholinguistic studies can be explained in terms of a categorical salience ranking.

# Pronoun Reference Resolution

Consider two approaches to salience ranking:

- categorical salience: referents are salient or non-salient
- gradient salience: referents are salient to varying degrees

Most psycholinguistic studies can be explained in terms of a categorical salience ranking.

- Order-of-mention studies (Gernsbacher and Hargreaves, 1988) show only advantage of first-mention.

# Pronoun Reference Resolution

Consider two approaches to salience ranking:

- categorical salience: referents are salient or non-salient
- gradient salience: referents are salient to varying degrees

Most psycholinguistic studies can be explained in terms of a categorical salience ranking.

- Order-of-mention studies (Gernsbacher and Hargreaves, 1988) show only advantage of first-mention.
- Centering Theory (Grosz and Sidner, 1986; Grosz et al., 1995) studies (e.g., Hudson-D'Zmura and Tanenhaus, 1997) of forward-looking center ranking show only $subjects > others$.

# Pronoun Reference Resolution

Consider two approaches to salience ranking:

- categorical salience: referents are salient or non-salient
- gradient salience: referents are salient to varying degrees

Most psycholinguistic studies can be explained in terms of a categorical salience ranking.

- Order-of-mention studies (Gernsbacher and Hargreaves, 1988) show only advantage of first-mention.
- Centering Theory (Grosz and Sidner, 1986; Grosz et al., 1995) studies (e.g., Hudson-D'Zmura and Tanenhaus, 1997) of forward-looking center ranking show only $subjects > others$.

The limitation with these studies is that they've compared salient entities to non-salient entities. But what does the processor do when a pronoun is compatible only with non-salient entities?

# Computing Salience of Referents

Salience can be computed from a number of factors.

# Computing Salience of Referents

Salience can be computed from a number of factors.

- Recency (Lappin and Leass, 1994)

# Computing Salience of Referents

Salience can be computed from a number of factors.

- Recency (Lappin and Leass, 1994)
- Syntactic Prominence (Mathews and Chodorow, 1988; McKoon et al., 1993)

# Computing Salience of Referents

Salience can be computed from a number of factors.

- Recency (Lappin and Leass, 1994)
- Syntactic Prominence (Mathews and Chodorow, 1988; McKoon et al., 1993)
- Semantic Prominence (Rose, 2005)

# Computing Salience of Referents

Salience can be computed from a number of factors.

- Recency (Lappin and Leass, 1994)
- Syntactic Prominence (Mathews and Chodorow, 1988; McKoon et al., 1993)
- Semantic Prominence (Rose, 2005)
- Parallelism (Stevenson et al., 1995)

# Computing Salience of Referents

Salience can be computed from a number of factors.

- Recency (Lappin and Leass, 1994)
- Syntactic Prominence (Mathews and Chodorow, 1988; McKoon et al., 1993)
- Semantic Prominence (Rose, 2005)
- Parallelism (Stevenson et al., 1995)
- Coherence Relations (Kehler, 2002; Stevenson et al., 2000)

# Computing Salience of Referents: Syntactic Prominence

Many psycholinguistic studies (e.g., Hudson-D'Zmura and Tanenhaus, 1997; Rose, 2005) assume a syntactic prominence hierarchy based on grammatical roles.

$$subject > object > others$$

# Computing Salience of Referents: Syntactic Prominence

Many psycholinguistic studies (e.g., Hudson-D'Zmura and Tanenhaus, 1997; Rose, 2005) assume a syntactic prominence hierarchy based on grammatical roles.

$$subject > object > others$$

Useful to explain such things as subject-preference in pronoun resolution.

# Computing Salience of Referents: Syntactic Prominence

Many psycholinguistic studies (e.g., Hudson-D'Zmura and Tanenhaus, 1997; Rose, 2005) assume a syntactic prominence hierarchy based on grammatical roles.

$$subject > object > others$$

Useful to explain such things as subject-preference in pronoun resolution.

But, ...

# Computing Salience of Referents: Semantic Prominence

In English, syntactic and semantic information are often conflated.

| John | hit | Matt. |
|---|---|---|
| SUBJECT | | OBJECT |
| AGENT | | PATIENT |

# Computing Salience of Referents: Semantic Prominence

In English, syntactic and semantic information are often conflated.

| John | hit | Matt. |
|---|---|---|
| SUBJECT | | OBJECT |
| AGENT | | PATIENT |

Thus, perhaps Syntactic Prominence effects are better explained in terms of Semantic Prominence effects.

# Computing Salience of Referents: Semantic Prominence

In English, syntactic and semantic information are often conflated.

| John | hit | Matt. |
|---|---|---|
| SUBJECT | | OBJECT |
| AGENT | | PATIENT |

Thus, perhaps Syntactic Prominence effects are better explained in terms of Semantic Prominence effects.

In Rose (2005), I observed in both psycholinguistic and corpus linguistic investigations that syntactic and semantic prominence contribute somewhat independently to salience ranking.

# Computing Salience of Referents: Semantic Prominence

In English, syntactic and semantic information are often conflated.

| John | hit | Matt. |
|---|---|---|
| SUBJECT | | OBJECT |
| AGENT | | PATIENT |

Thus, perhaps Syntactic Prominence effects are better explained in terms of Semantic Prominence effects.

In Rose (2005), I observed in both psycholinguistic and corpus linguistic investigations that syntactic and semantic prominence contribute somewhat independently to salience ranking.

1. John$_i^{\star\bullet}$ could easily hit Matt$_j$. He$_{i>j}$ ...

# Computing Salience of Referents: Semantic Prominence

In English, syntactic and semantic information are often conflated.

| John | hit | Matt. |
|---|---|---|
| SUBJECT | | OBJECT |
| AGENT | | PATIENT |

Thus, perhaps Syntactic Prominence effects are better explained in terms of Semantic Prominence effects.

In Rose (2005), I observed in both psycholinguistic and corpus linguistic investigations that syntactic and semantic prominence contribute somewhat independently to salience ranking.

1. John$_i^{\star\bullet}$ could easily hit Matt$_j$. He$_{i>j}$ ...
2. Matt$_j^{\star}$ was easy for John$_i^{\bullet}$ to hit $\emptyset_j$. He$_{i=j}$ ...

# Computing Salience of Referents: Coherence Relations

The logical relationship between adjacent utterances may influence pronoun resolution preferences (Kehler, 2002).

# Computing Salience of Referents: Coherence Relations

The logical relationship between adjacent utterances may influence pronoun resolution preferences (Kehler, 2002).

1. Luke hit Max and then he$_{i/\#j}$ ran home. NARRATIVE

# Computing Salience of Referents: Coherence Relations

The logical relationship between adjacent utterances may influence pronoun resolution preferences (Kehler, 2002).

1. Luke hit Max and then he$_{i/\#j}$ ran home. NARRATIVE
2. Luke hit Max because he$_{\#i/j}$ ran home. CAUSAL

# Computing Salience of Referents: Coherence Relations

The logical relationship between adjacent utterances may influence pronoun resolution preferences (Kehler, 2002).

1. Luke hit Max and then he$_{i/\#j}$ ran home. NARRATIVE
2. Luke hit Max because he$_{\#i/j}$ ran home. CAUSAL

Stevenson et al. (1994, 2000) account:

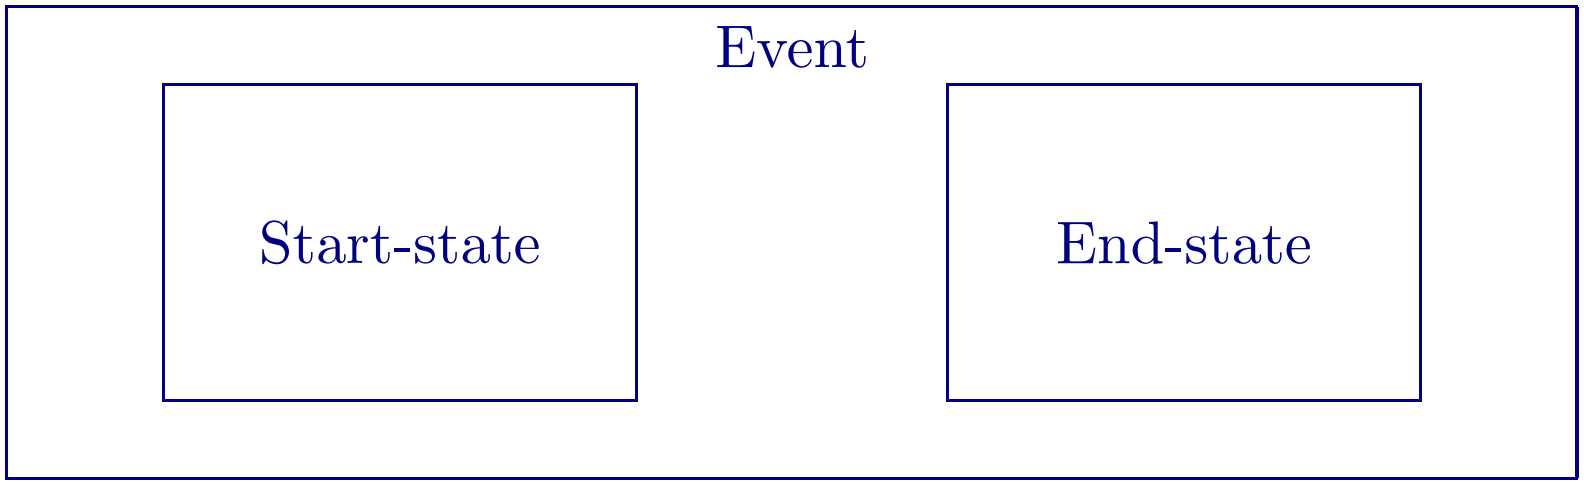

# Computing Salience of Referents: Coherence Relations

The logical relationship between adjacent utterances may influence pronoun resolution preferences (Kehler, 2002).

1. Luke hit Max and then he$_{i/\#j}$ ran home. NARRATIVE
2. Luke hit Max because he$_{\#i/j}$ ran home. CAUSAL

Stevenson et al. (1994, 2000) account:

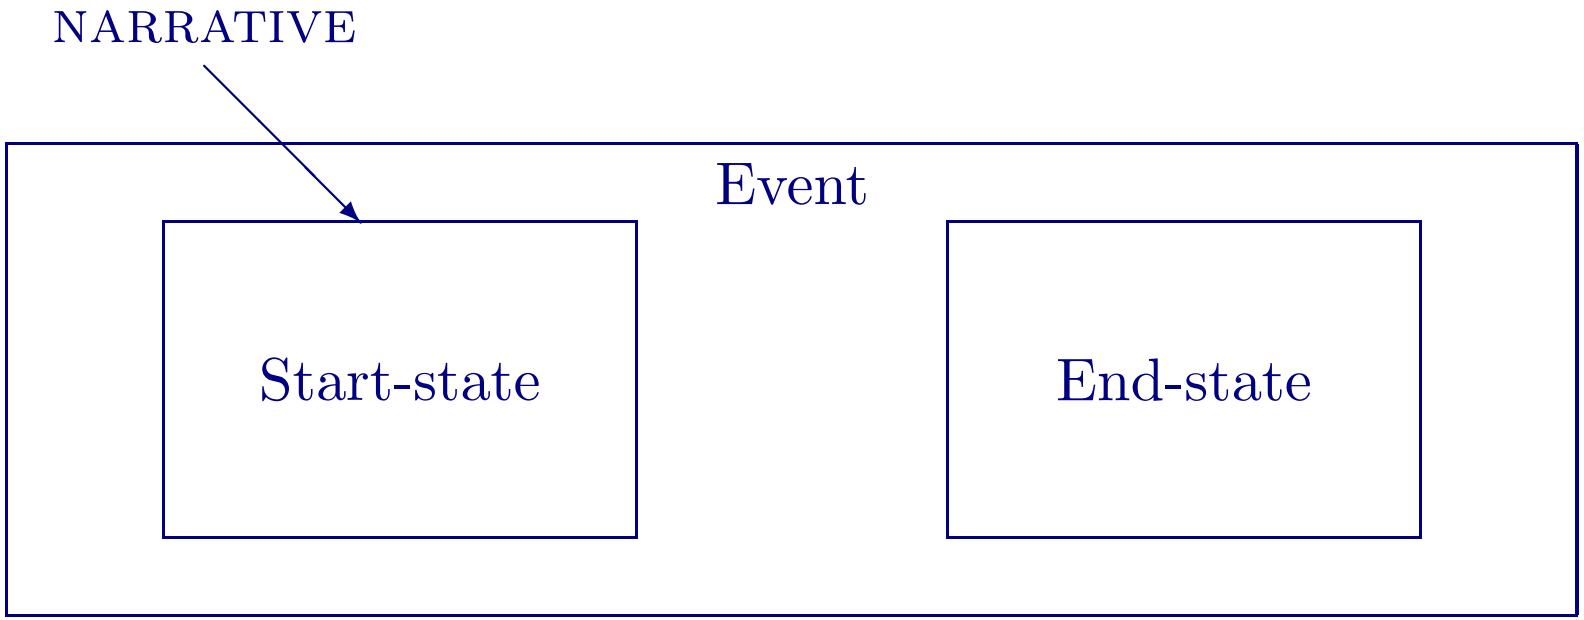

# Computing Salience of Referents: Coherence Relations

The logical relationship between adjacent utterances may influence pronoun resolution preferences (Kehler, 2002).

1. Luke hit Max and then he$_{i/\#j}$ ran home. NARRATIVE
2. Luke hit Max because he$_{\#i/j}$ ran home. CAUSAL

Stevenson et al. (1994, 2000) account:

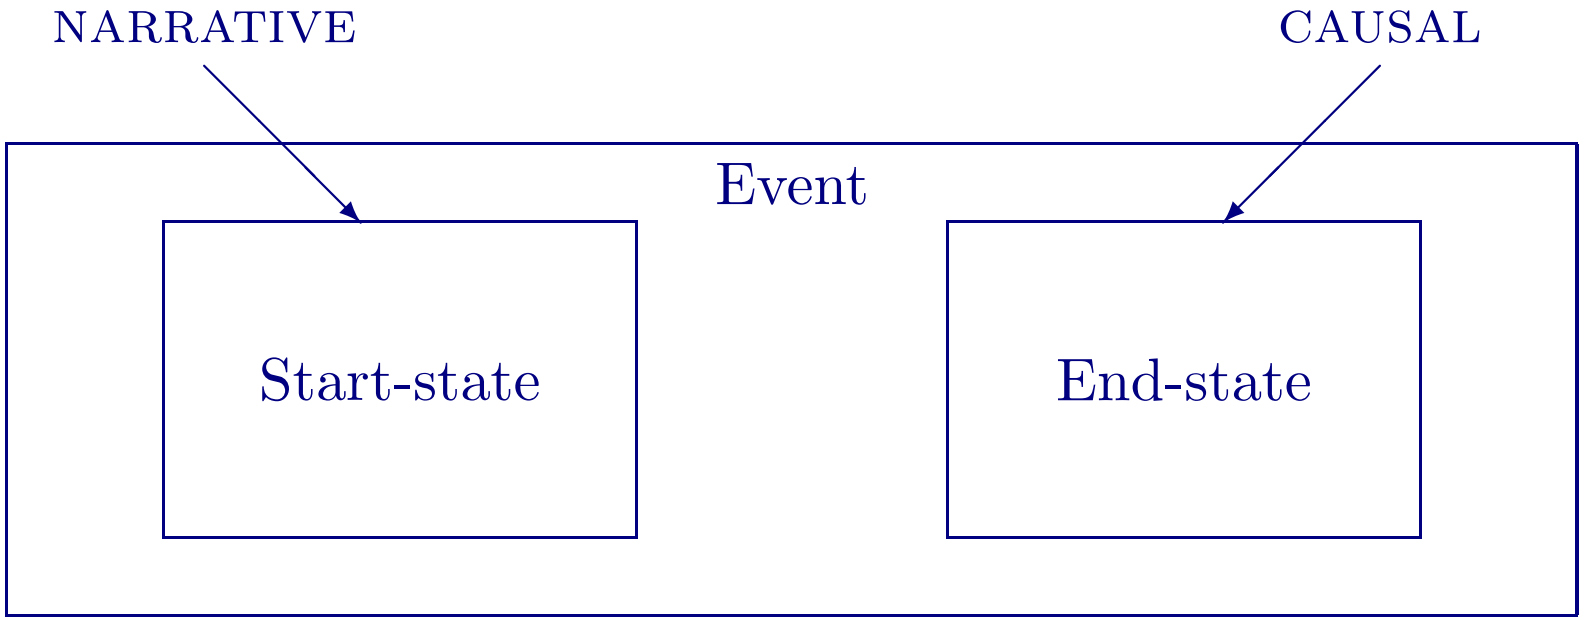

# Experiment

There are three questions addressed by this study.

# Experiment

There are three questions addressed by this study.

- Primary: Is salience categorical or gradient?
- Secondary: How do syntactic and semantic prominence contribute to salience?
- Secondary: What influence do coherence relations across utterances have on resolution preferences with non-salient entities?

# Experiment

There are three questions addressed by this study.

- Primary: Is salience categorical or gradient?
- Secondary: How do syntactic and semantic prominence contribute to salience?
- Secondary: What influence do coherence relations across utterances have on resolution preferences with non-salient entities?

# Experiment

There are three questions addressed by this study.

- Primary: Is salience categorical or gradient?
- Secondary: How do syntactic and semantic prominence contribute to salience?
- Secondary: What influence do coherence relations across utterances have on resolution preferences with non-salient entities?

# Experiment

In Rose (2005), I used spray/load constructions to compare syntactic and semantic prominence effects.

| John | sprayed | some paint | on | a wall. | It ... |
|---|---|---|---|---|---|
| AGENT | | THEME | | LOCATION | |

| John | sprayed | a wall | with | some paint. | It ... |
|---|---|---|---|---|---|
| AGENT | | LOCATION | | THEME | |

# Experiment

In Rose (2005), I used spray/load constructions to compare syntactic and semantic prominence effects.

| John | sprayed | some paint | on | a wall. | It ... |
|---|---|---|---|---|---|
| AGENT | | THEME | | LOCATION | |

| John | sprayed | a wall | with | some paint. | It ... |
|---|---|---|---|---|---|
| AGENT | | LOCATION | | THEME | |

This is useful to test categorical vs. gradient salience: John (subject, agent, human, named entity) is most salient. Thus, pronominal reference with third-person singular pronoun permits test of reference to non-salient entities.

# Experiment

Two experimental factors:

# Experiment

Two experimental factors:

- ORDER: theme-location, location-theme

# Experiment

Two experimental factors:

- ORDER: theme-location, location-theme
- RELATION: narrative (*and then*), causal (*because*)

# Experiment

Two experimental factors:

- ORDER: theme-location, location-theme
- RELATION: narrative (*and then*), causal (*because*)

Thus, four possible conditions:

1. John sprayed some paint on a wall and then it ...
   (theme-location, narrative)
2. John sprayed a wall with some paint and then it ...
   (location-theme, narrative)
3. John sprayed some paint on a wall because it ...
   (theme-location, causal)
4. John sprayed a wall with some paint because it ...
   (location-theme, causal)

# Experiment

- 24 test items pseudo-randomized and combined with 101 items from an unrelated experiment to make total of 125 items.

# Experiment

- 24 test items pseudo-randomized and combined with 101 items from an unrelated experiment to make total of 125 items.
- 36 native English undergraduate students participated in the experiment.

# Experiment

- 24 test items pseudo-randomized and combined with 101 items from an unrelated experiment to make total of 125 items.
- 36 native English undergraduate students participated in the experiment.
- Participants were asked to write what they regarded to be a natural completion for each item.

# Results and Analysis

Completions were read and marked for the interpretation of the pronoun.

641 completions unambiguously referred to either the theme or location.

# Results and Analysis

Completions were read and marked for the interpretation of the pronoun. 641 completions unambiguously referred to either the theme or location.

Results are presented from two perspectives:

- Object-preference: number of choices for object minus number of choices for oblique
- Theme-preference: number of choices for theme minus number of choices for location

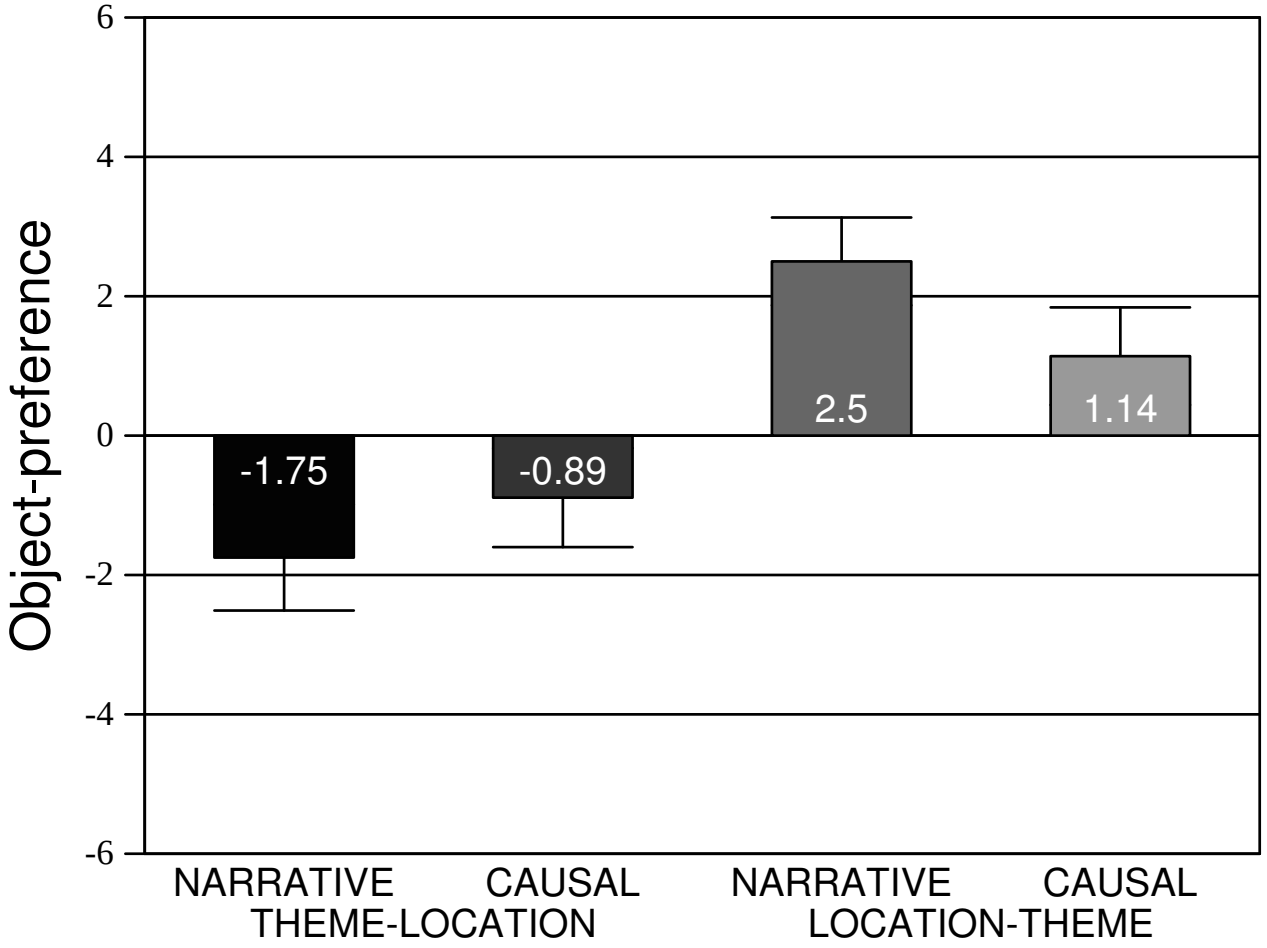


| | by subjects | | by items | |
|---|---|---|---|---|
| ORDER | $F(1,35) = 58.2$ | $p < 0.001$ | $F(1,23) = 18.3$ | $p < 0.001$ |
| RELATION | $F(1,35) < 1.0$ | n.s. | $F(1,23) < 1.0$ | n.s. |
| ORDER*RELATION | $F(1,35) = 8.4$ | $p < 0.01$ | $F(1,23) = 3.2$ | $p = 0.085$ |

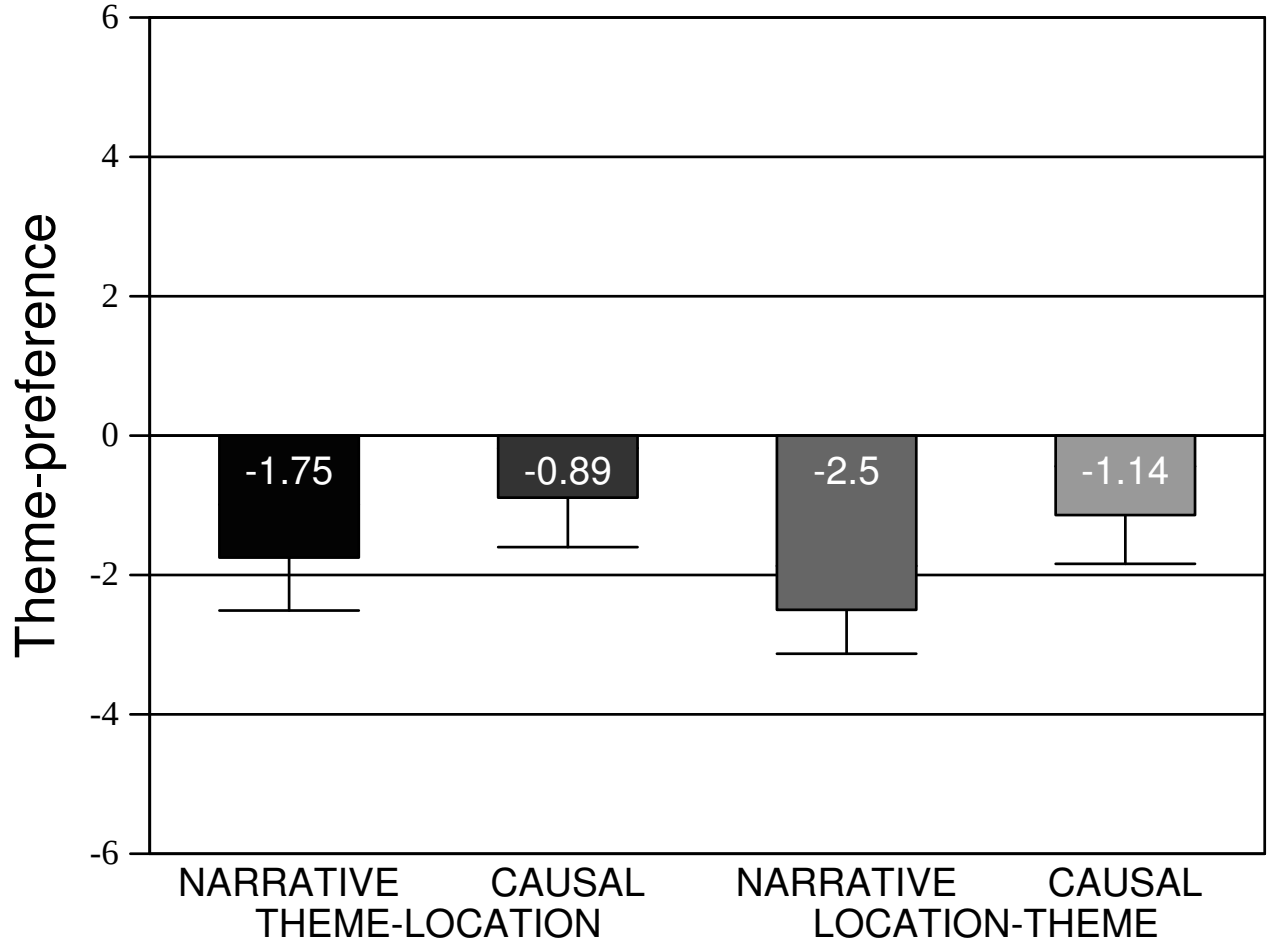


| | by subjects | | by items | |
|---|---|---|---|---|
| ORDER | $F(1,35) = 2.8$ | $p = 0.10$ | $F(1,23) = 2.2$ | $p = 0.15$ |
| RELATION | $F(1,35) = 8.4$ | $p < 0.01$ | $F(1,23) = 3.2$ | $p = 0.085$ |
| ORDER*RELATION | $F(1,35) < 1.0$ | n.s. | $F(1,23) < 1.0$ | n.s. |

# Results and Analysis

- Q: Is salience better represented as a categorical or gradient measure?

  A: gradient measure

# Results and Analysis

- Q: Is salience better represented as a categorical or gradient measure?
  A: gradient measure
- Q: Which determines the salience of entities for pronoun resolution: syntactic prominence or semantic prominence?
  A: semantic prominence (in this experiment)

# Results and Analysis

- Q: Is salience better represented as a categorical or gradient measure?
  A: gradient measure
- Q: Which determines the salience of entities for pronoun resolution: syntactic prominence or semantic prominence?
  A: semantic prominence (in this experiment)
- Q: What effect do coherence relations have on pronoun resolution preferences?
  A: no meaningful effect (i.e., no flip-flop in resolution preferences)

# Discussion

## Categorical vs. Gradient Salience

- Results provide some empirical foundation for psycholinguistic models which use gradient salience (e.g., Gernsbacher and Hargreaves, 1988; Hudson-D'Zmura and Tanenhaus, 1997).
- Results provide psycholinguistic footing for computational implementations of which use a gradient salience ranking (e.g., Kennedy and Boguraev, 1996; Lappin and Leass, 1994).

# Discussion

## Categorical vs. Gradient Salience

- Results provide some empirical foundation for psycholinguistic models which use gradient salience (e.g., Gernsbacher and Hargreaves, 1988; Hudson-D'Zmura and Tanenhaus, 1997).
- Results provide psycholinguistic footing for computational implementations of which use a gradient salience ranking (e.g., Kennedy and Boguraev, 1996; Lappin and Leass, 1994).

## Syntactic vs. Semantic Prominence

- Results are at some variance with previous work (Rose, 2005)
- Semantic prominence is a significant factor in computing salience.

# Further Work

Salience is gradient, but *how* gradient is it?

- Is salience a discrete or a continuous scale?
- What bounds does it have? if any?

# Further Work

Salience is gradient, but *how* gradient is it?

- Is salience a discrete or a continuous scale?
- What bounds does it have? if any?

Is salience gradient cross-linguistically?

- In English, salience is gradient.
- Could there be languages in which salience is categorical?

Thank You!

# References


Almor, A. (1999). Noun-phrase anaphora and focus: The informational load hypothesis. *Psychological Review*, 106:748–765.

Gernsbacher, M. A. and Hargreaves, D. (1988). Accessing sentence participants: The advantage of first mention. *Journal of Memory and Language*, 27:699–717.

Giora, R. (2003). *On Our Mind: Salience, Context, and Figurative Language*. Oxford University Press, Oxford, UK.

Greene, S., McKoon, G., and Ratcliff, R. (1992). Pronoun resolution and discourse models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18:266–283.

Grosz, B., Joshi, A., and Weinstein, S. (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21:203–225.

Grosz, B. and Sidner, C. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12:175–204.

Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst.

Hirst, G. (1981). *Anaphora in Natural Language Understanding: A Survey*. Springer-Verlag, Berlin.

Hudson-D'Zmura, S. and Tanenhaus, M. (1997). Assigning antecedents to ambiguous pronouns: The role of the center of attention as the default assignment. In Walker, M., Joshi, A., and Prince, E., editors, *Centering Theory in Discourse*, pages 199–226. Clarendon Press, Oxford.

Karttunen, L. (1976). Discourse referents. In McCawley, J., editor, *Syntax and Semantics, Vol. 7: Notes from the Linguistic Underground*, pages 363–385. Academic Press, New York.

Kehler, A. (2002). *Coherence, Reference, and the Theory of Grammar*. CSLI Publications, Stanford University, CA.

Kennedy, C. and Boguraev, B. (1996). Anaphora for everyone: Pronominal anaphora resolution without a parser. In *Proceedings of the 16th International Conference on Computational Linguistics (COLING '96)*, pages 113–118, Copenhagen, Denmark.

Lappin, S. and Leass, H. (1994). An algorithm for pronominal anaphora resolution. *Computational Linguistics*, 20:535–561.

Mathews, A. and Chodorow, M. (1988). Pronoun resolution in two-clause sentences: Effects of ambiguity, antecedent location, and depth of embedding. *Journal of Memory and Language*, 27:245–260.

McKoon, G., Ratcliff, R., Ward, G., and Sproat, R. (1993). Syntactic prominence effects on discourse processes. *Journal of Memory and Language*, 32:593–607.

Mitkov, R. (2002). *Anaphora Resolution*. Longman, London.

Rose, R. (2005). *The Relative Contribution of Syntactic and Semantic Prominence to the Salience of Discourse Entities*. PhD thesis, Northwestern University.

Stevenson, R., Crawley, R., and Kleinman, D. (1994). Thematic roles, focus and the representation of events. *Language and Cognitive Processes*, 9:519–548.

Stevenson, R., Knott, A., Oberlander, J., and McDonald, S. (2000). Interpreting pronouns and connectives: Interactions among focusing, thematic roles and coherence relations. *Language and Cognitive Processes*, 15(3):225–262.

Stevenson, R., Nelson, A., and Stenning, K. (1995). The role of parallelism in strategies of pronoun comprehension. *Language and Speech*, 38(4):393–418.

Tetreault, J. (2001). A corpus-based evaluation of centering theory and pronoun resolution. *Computational Linguistics*, 27:507–520.