Accenting and Information Status

Julia Hirschberg

CS 4706

Information Status

- Topic/comment, theme/rheme

  The orangutan we wanted to buy escaped from the pet store.
- Focus of attention

  I only bought ***candy*** for that orangutan.
- Given/new

  I only bought candy for that orangutan. I would never buy an ape drugs!
- All commonly signaled in human speech by intonation

Today: Acent and Given/New

- Motivation in speech technology
- Models of Given/New
- Experiments on Given/New and pitch accent
- Possible models of intonation wrt given/new entities
- How might we identify given/new information automatically?
- How should we produce given/new information appropriately?
- Why is this important?

A Simple Definition

- Given: ***Recoverable*** from some form of context or, what a Speaker believes to be in a Hearer's consciousness
- New: ***Not*** recoverable from context or, what a Speaker believes is ***not*** in a Hearer's consciousness

Role in Speech Technologies

- TTS: Natural production
  - *Given information* is *often* ***deaccented***
  - *New information* is *usually* ***accented***
- ASR: Improved recognition
  - *Given information* may already have been recognized earlier
  - *New information* may be important cue to topic shift
- Summarization: Improved precision
  - *Given information* less likely to be included in a summary; *new information* more likely

- Spoken Dialogue Systems: Grounding
  - Critical for system to convey what is given and what is new to facilitate Hearer comprehension

Prince ’81: A More Complex Model

- Speaker (S) and Hearer (H), in a discourse, construct a discourse model
  - Includes discourse entities, attributes, and links between entities
  - Discourse entities: individuals, classes, exemplars, substances, concepts (NPs)
- Entities when first introduced are new
  - *Brand-new* (H must create a new entity)

    My dog bit a rhinoceros this morning.

- *Unused* (H already knows of this entity)

  The sun came out this morning.

- Evoked entities are old, or 'given' -- already in the discourse
  - *Explicitly evoked (in text or speech)*

    The rhinoceros was wearing suspenders. Rather unusual for a rhino.
  - *Situationally evoked*

    Watch out for the snake!
- Inferables are also old, or 'given'

  I bought a new car. The gear shift is a bit tricky.

Prince '92: A Still More Complex Model

- Hearer-centric information status:
  - Given: what S believes H has in his/her consciousness
  - New: what S believes H does ***not*** have in his/her consciousness
- But discourse entities may also be *given* and *new* wrt the current discourse
  - Discourse-old: already evoked in the discourse
  - Discourse-new: not evoked

The ***stars*** are very bright tonight (Hearer-given; Discourse-new)

When I see ***stars*** this bright, I think of my vacations in the mountains. (Hearer-given; Discourse-given)

My friend ***Buddy*** and I would sneak out late at night. (Hearer-new; Discourse-new)

I said, “My friend ***BUDDY***…” (Hearer-new; Discourse-given)

Given/New and Pitch Accent

- New information is often accented and given information is often deaccented (Halliday ‘67, Brown ‘83, Terken ‘84)
  - But there are many exceptions: a simple TTS rule: accent ‘new’ and deaccent ‘given’ will make 25-30% errors
  - How can we reduce these errors, to produce human-like intonation?

Brown ‘83: Accent Status and Subclasses of Given/New

- Speech elicitation in laboratory
  - 12 Scottish-English undergrads
  - *A* describes a diagram for *B* to draw, which *B* cannot see

    Draw a black triangle.

    Draw a circle in the middle.

    Draw a blue triangle next to the black one with a line from the top angle to the bottom.
- Analysis: based on Prince ‘81 categories with modifications

  - Brand-new (a triangle), given:inferrable (middle, angle), given:contextually evoked (the page), given:'textually' evoked (divided into current topic vs. earlier mention)
  - Accent status of all entity-referring NPs
- Results:
  - Brand-new information accented (87%)
    - Note: new entity/old expression issue
  - Given: contextually evoked information deaccented (98%)
  - Given: 'textually' evoked deaccented (current topic 100%; earlier: 96%)
  - Given: inferable information ***accented*** (79%)

Boston Directions Corpus (Hirschberg & Nakatani '96)

- Experimental Design
  - 12 speakers: 4 used
  - Spontaneous and read versions of 9 direction-giving tasks (monologues)
- Corpus: 50m read; 67m spon
- Labeling
  - Prosodic: ToBI intonational labeling
  - Given/new (Prince '92), grammatical function, p.o.s.,…

Boston Directions Corpus: Describe how to get to MIT from Harvard 

d1: dsp1: step 1: enter and get token
first
enter the Harvard Square T stop
and buy a token

d2: dsp2: inbound on red line
then
proceed to get on the
inbound
um
Red Line
uh subway

dp3 dsp3: take subway from hs, to cs to ks
and
take the subway
from Harvard Square
to Central Square
and then to Kendall Square

dp4: dsp4: get off T.
then get off the T

Hearer and Discourse Given/New Labeling

first
enter the Harvard Square T stop
and buy a token
then
proceed to get on the
inbound
um
Red Line
uh subway
and
take the subway
from Harvard Square
to Central Square
and then to Kendall Square
then get off the T

Hearer and Discourse Given/New Labeling

first
enter <HG/DN the Harvard Square T stop>
and buy <HI/DN a token>
then
proceed to get on <HI/DN the
inbound
um
Red Line
uh subway>
and
take <HG/DG the subway>
from <HG/DG Harvard Square>
to <HG/DN Central Square>
and then to <HG/DN Kendall Square>
then get off <HG/DG the T>

Does Given/New Status Predict Deaccenting?

NPa	HG	HI	HN	DG	DN
Deaccented	37.1%	53.9%	26.2%	43.3%	38.8%
Total	1009	406	130	596	950

HG: Hearer Given HI: Hearer Inferable HN: Hearer New DG: Discourse Given DN: Discourse New

39.4% of (H or D) Given items deaccented…

36.9% of (H or D) New Items are deaccented…

And….Bard'99: Givenness, deaccenting and intelligibility

- Speech elicited in laboratory
  - Glasgow Scottish-English Map Task
    - Each has a slightly different map
    - *A* traces a route described by *B*
- Analysis
  - Compare repeated mentions of same items (i.e. given items) wrt accent status
    - Within dialogue
    - Across dialogue
- Findings

  - Deaccenting rare in repeated mentions (within 15% and across 6% dialogues)
  - But repeated mentions were `less intelligible'
- Caveats:
  - Were they really identifying 'deaccenting' (the absence of a pitch accent)?
  - Were mentions within speaker or across speaker?
  - Some more questions to ask….

What else is going on?

- Given/new and grammatical function
- Hypothesis: ***how*** discourse entities are evoked in a discourse influences accent status
- E.g., How might grammatical function and surface position interact with the accentuation of 'given' items?
- Cases:
  - X has not been mentioned in the prior context
  - X has been mentioned, with the same grammatical function/surface position
  - X has been mentioned but with a different grammatical function/surface position

Experimental Design

- Major problem:
  - How to elicit ‘spontaneous’ productions while varying desired phenomena systematically?
  - Key: simple variations and actions can capitalize upon natural tendency to associate grammatical functions with particular thematic roles for a given set of verbs

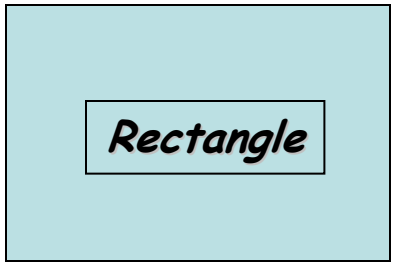
Rectangle

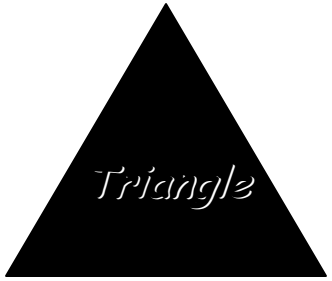
Triangle

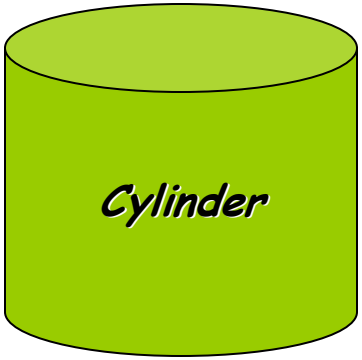
Cylinder

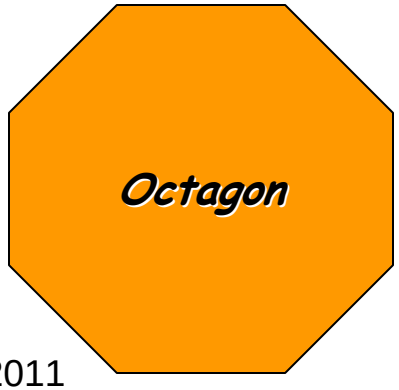
Octagon

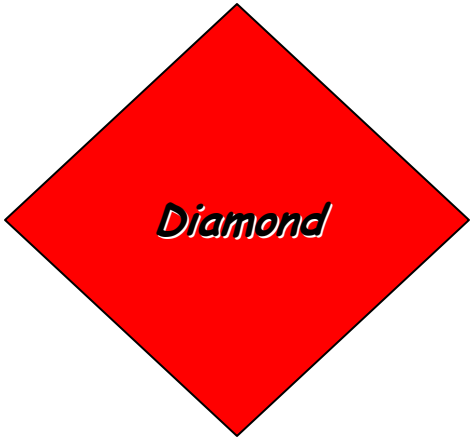
Diamond

Context 1

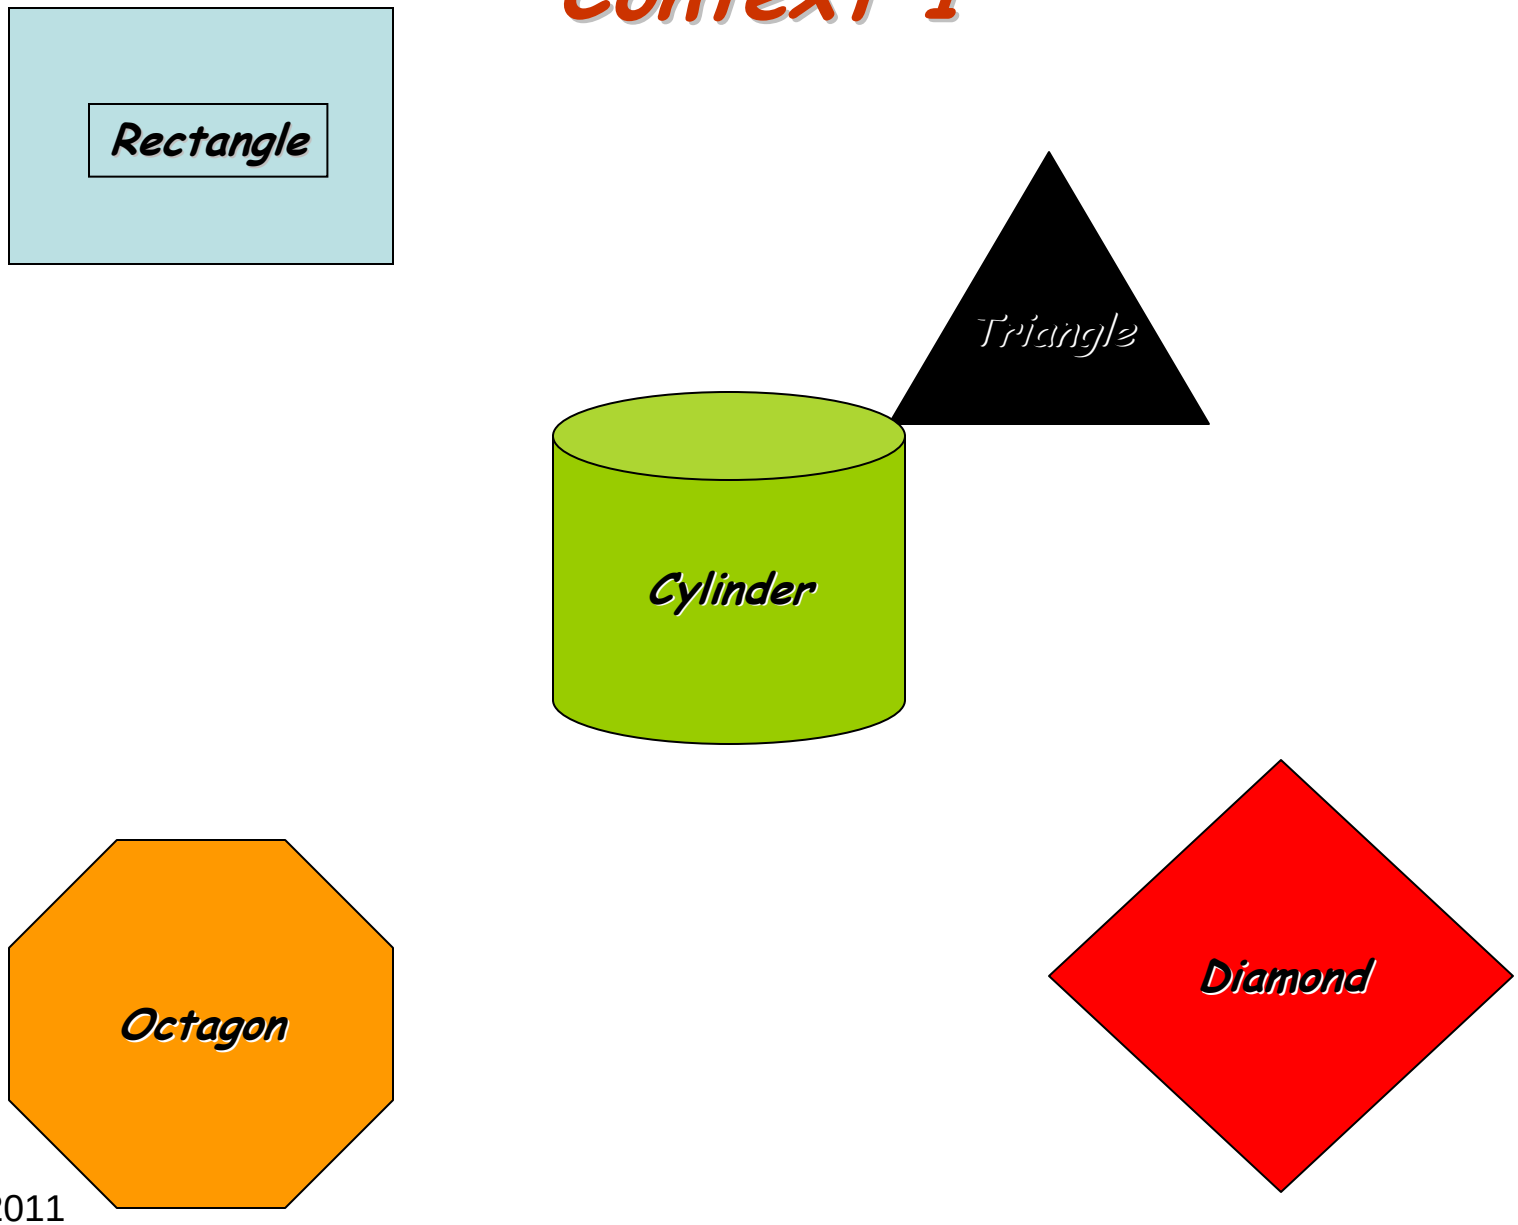

Context 2

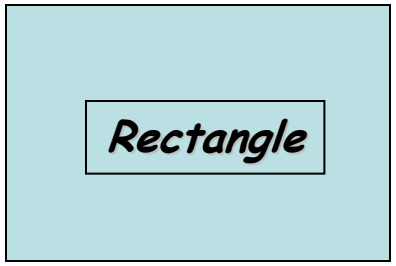

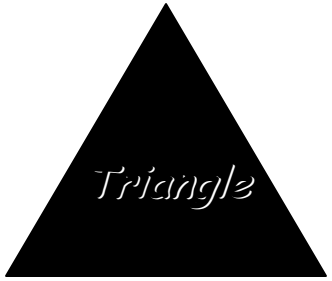

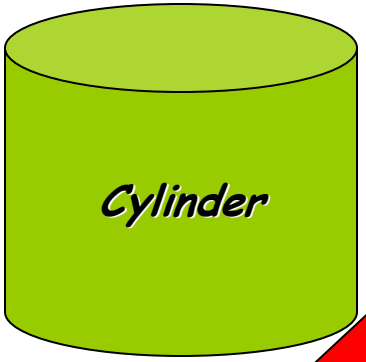

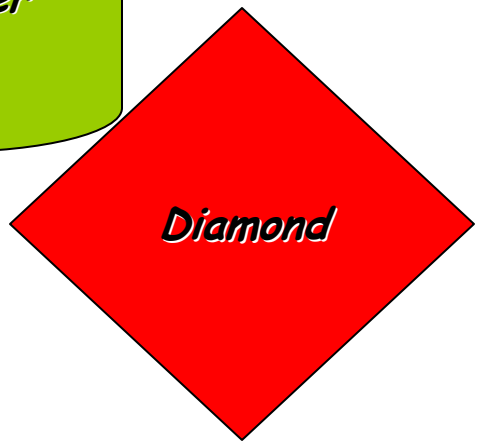

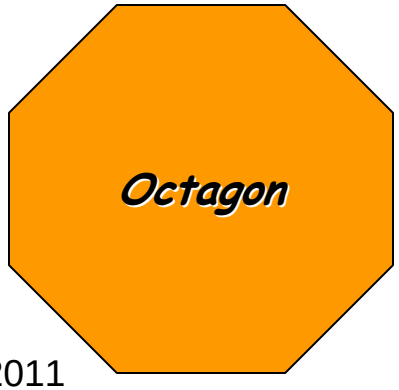

Context 3

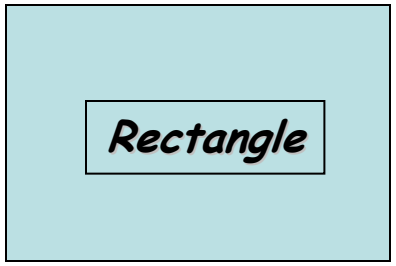

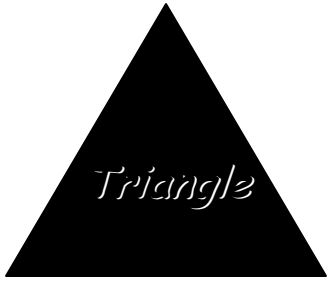

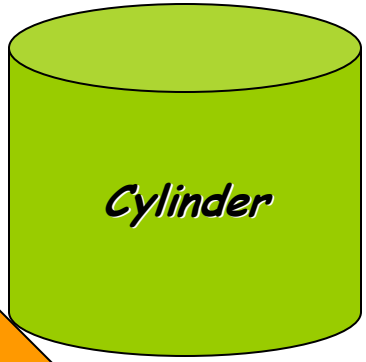

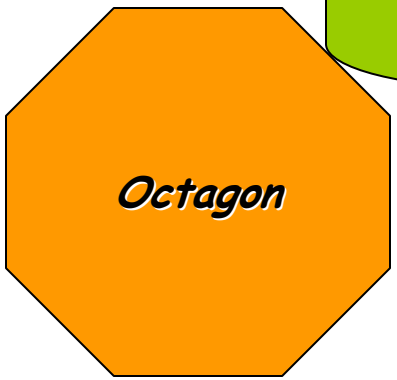

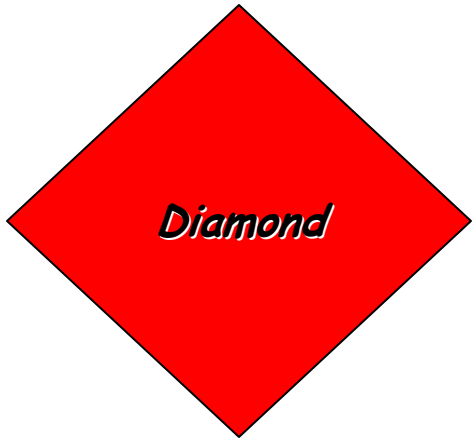

Target(A)

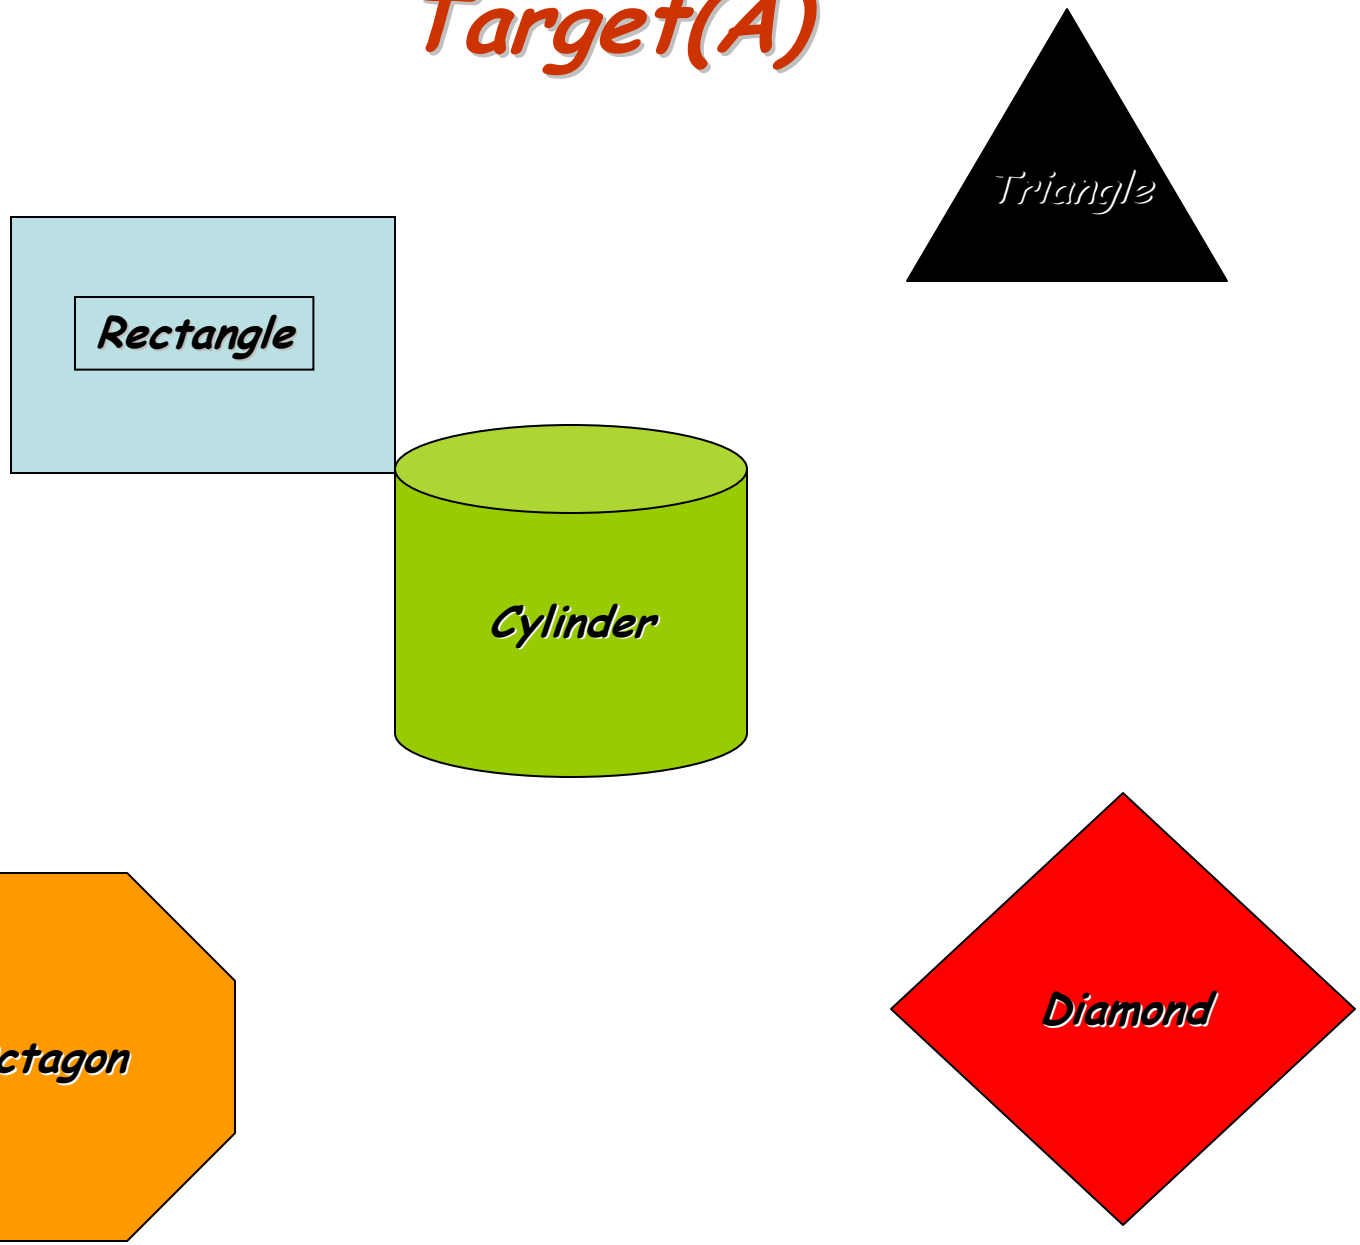

Target(B)

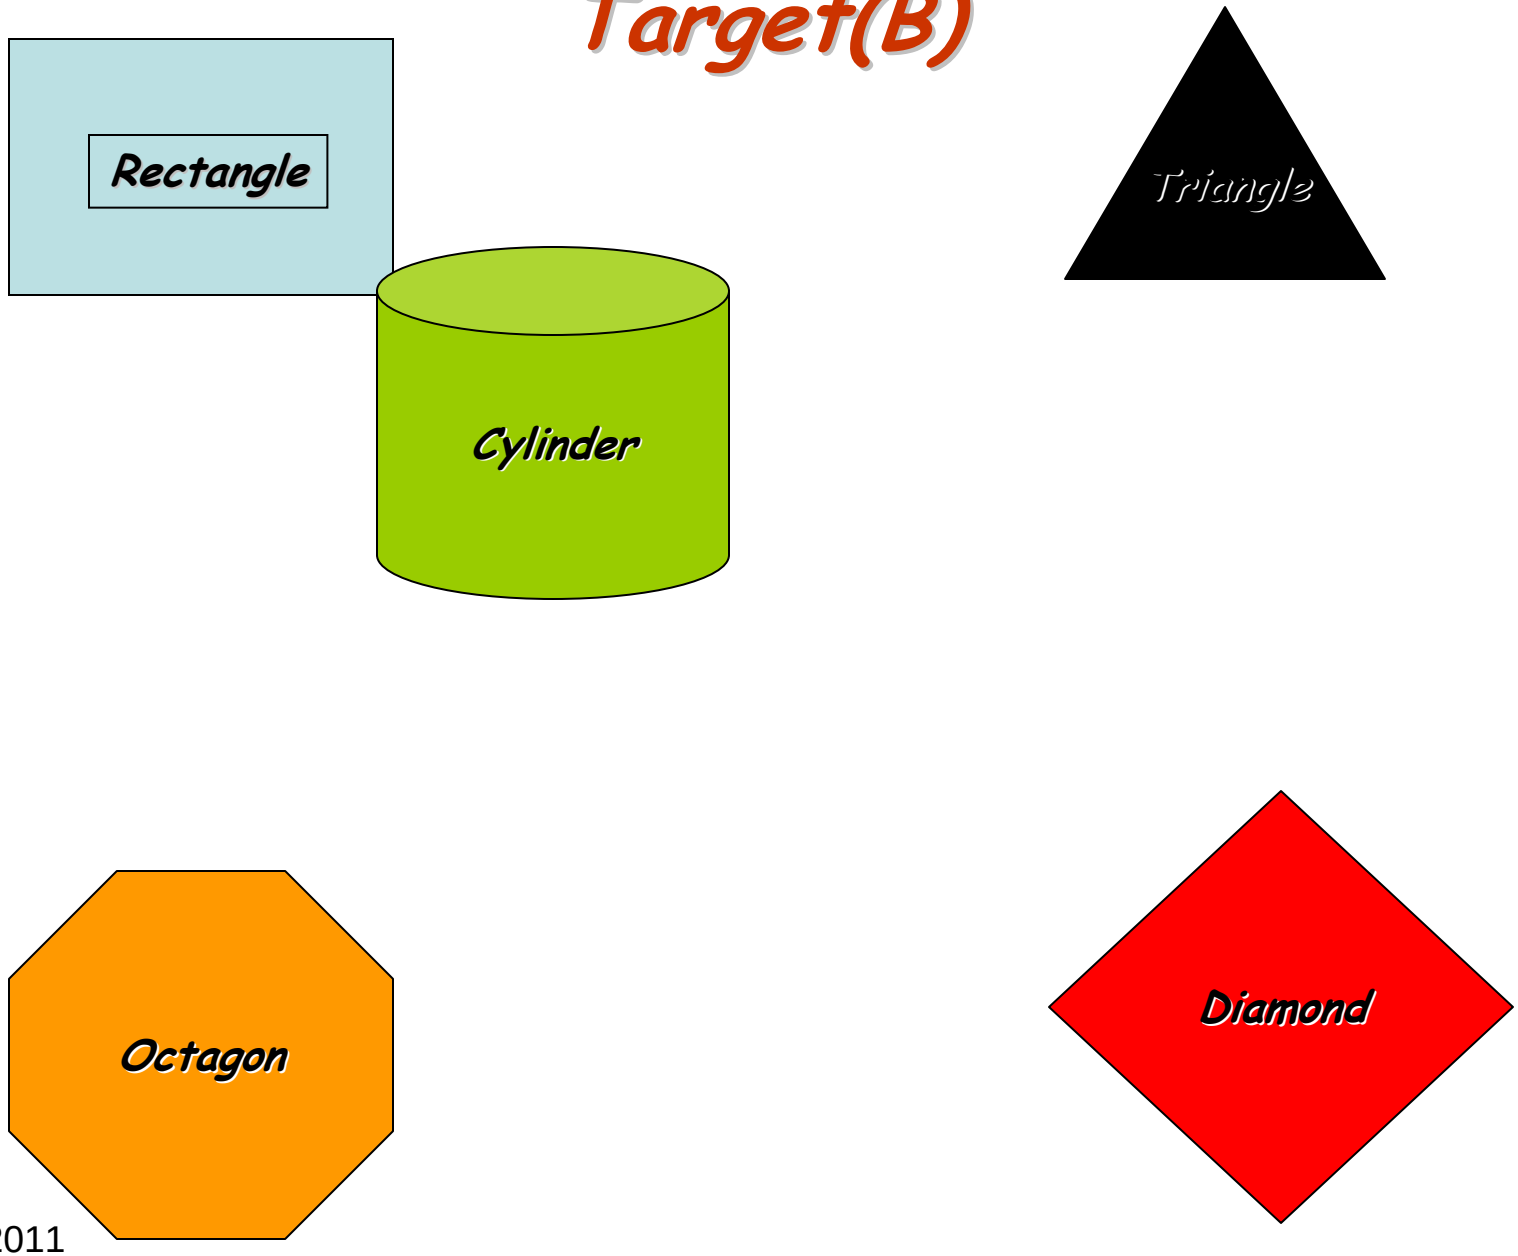

Materials

- 9 objects in visual display
- 3 event types:
  - X covers Y (subject, object)
  - X pushes Y against Z (subject, object, pp-object)
  - X touches Y (subject, object)
- 75 scenarios of 4 sequences of actions each
  - 3 context turns (all containing the same given item)
  - 1 target turn (always containing the same given item)
  - 3x3 design (given item is subj, direct object or pp-obj in context and same or not in target) with 5 scenarios per cell
  - 2 controls: all new, all given objects (15 scenarios each)
  - Presented in random order

Experimental Conditions

- 10 native speakers of standard American English
- Subject and experimenter in soundproof booth
- Subject told to describe scenes to confederate outside the booth, visible but with providing no feedback
- 10 practice scenarios
- ~20 minutes per subject

Prosodic Analysis

- Target turns excised and analyzed by two judges independently for location of pitch accents for each referring expression: accented (2), unsure (1), deaccented (0) → accentedness score from 0-4 (81% agreement for 0 and 2 scores)
- Accent scores calculated by adding the two judges scores for each item (range: 0-4)

Grammatical Role/Surface Position Accenting 'Score'

	CONTEXT		TARGET	
GIVEN		Subj	D-obj	Pp-obj
	Subj	2.1	3.6	3.2
	D-obj	3.3	0.6	1.6
	Pp-obj	3.0	1.4	0.7
NEW		3.7	3.8	--

Findings

- Items that ***differ from context to target in grammatical function or surface position*** tend to be accented
- Items that ***share grammatical function and surface position*** tend to be deaccented
- ***Subjects were accented more than objects***, even if previously mentioned in same role
- Direct objects and pp-objects differ more from subjects than from each other

Given/New Isn't Just About Discourse Entities

- Consider e.g. Subj→D.O. variation

  The TRIANGLE touches the CYLINDER.

  The triangle touches the DIAMOND.

  The triangle touches the OCTAGON.

  The RECTANGLE touches the TRIANGLE.

- An entity may be 'given' or 'new' ***wrt the role it plays in the discourse***

How can we determine automatically whether a discourse entity is given or new?

- A rule-based approach:
  - Stem the content words in the discourse
  - Select a window within which incoming items with the same stem as a previous entity and within this window will be labeled ‘given’
    - Other items are ‘new’
- Is this hearer-based? Discourse-based?
- How well does it predict pitch accent?
  - 65-75% accurate (precision) depending on genre, domain

What else can we do?

- Instead of just accenting new and deaccenting given items
  - Keep track of given ‘type’ (evoked or inferable)
  - Keep track of grammatical function of discourse entities when introduced (subject, direct object, pp-object)
- What else? downstepped contours and given/new

How important is it to accent given/new items appropriately?

- Are listeners sensitive to intonational correlates of information status?
- Evidence that **'appropriate' accentuation facilitates comprehension**:
  - Birch & Clifton (1995): appropriateness speeds makes-sense judgments in Q&A pairs
  - Bock & Mazzella (1983): comprehension time of denial-counterassertion pairs
  - Davidson (2001): phoneme-monitoring in denial-counterassertion pairs
  - Terken & Nooteboom (1987), Nooteboom & Kruyt (1987): verification latencies of target words

Intonational cues in on-line processing

- Dahan et al. (2002): accentuation effects referential interpretation even at very early stages of processing
- Used eye-tracking to monitor listeners' fixations on pictured entities as they heard instructions to manipulate these entities on computer screen
- Examined moment-by-moment recognition of accented vs. unaccented words which share a primary-stressed initial syllable (e.g. candy/candle).

Dahan et al. (2002)

- Example discourse:

  "Put the CANDLE below the triangle.

  Now put the CAN | DLE above the square."

utt 1	utt 2	pred. fixation
<candle>	Now put the [kæn] ...	= candle
<candle>	Now put the [**kæn**] ...	= candy
<candy>	Now put the [kæn] ...	= candy
<candy>	Now put the [**kæn**] ...	= candle

candle-CANDLE condition

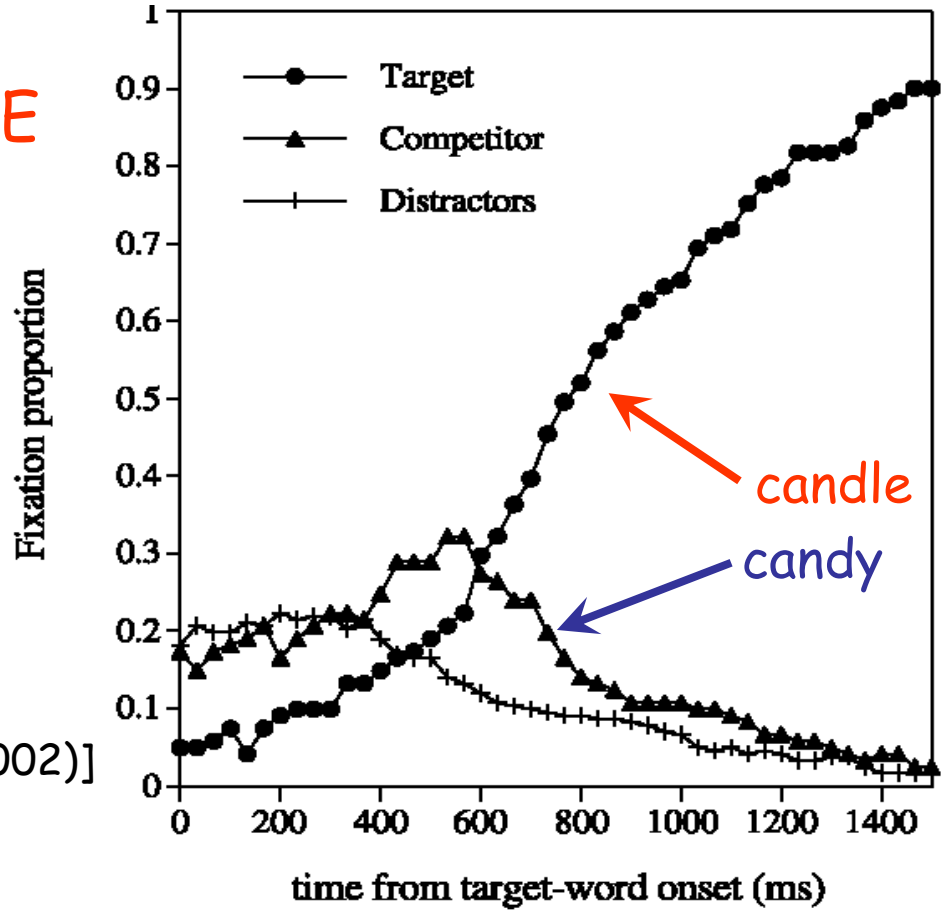

[From Dahan et al. (2002)]

- Competition from “candy” upon hearing accented [**kæn**].
- Accentuation is used by listeners to process discourse representations on-line, as a word is unfolding.

Accent, Given/New, and Grammatical Function

- Some more evidence about grammatical function….
- Dahan et al '02 also examine conditions in which the antecedent and target did not share grammatical role:

  "Put the necklace below the candle.

  Now put the CANDLE above the square."

  – NO competitor effects (i.e. looks to "candy" upon hearing accented [**kæn**])

- "Put the necklace below the candle.
  Now put the CANDLE above the square."
- Prediction: "candle" is given ⇒ competition from "candy" upon hearing accented [**kæn**].

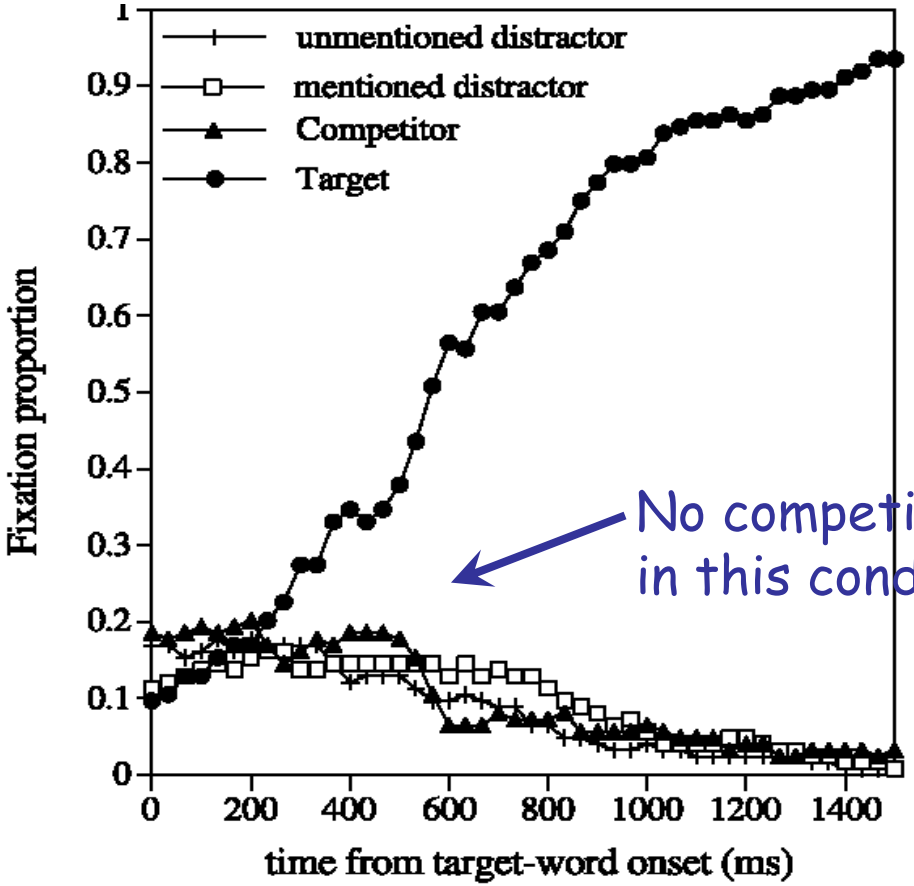

[From Dahan et al. (2002)]

No competitor effects in this condition!

Here, accent serves to cue shift in "focus".

[See also Terken & Hirschberg (1994)]

Grammatical Role or Syntactic parallelism?

- Dahan et al.'s and Terken & Hirschberg's data confound grammatical role with syntactic parallelism
- Venditti et al '02,'03: syntactic parallelism (NOT just persistence of grammatical role) affects interpretation of nuclear-accented pronouns

John hit Bill and then HE ... hit George. (N2 pref)

... ran away. (less N2 pref)

- N2 pref when syntatically parallel clauses

Next Class

- Back end synthesis and TTS evaluation