2006-01-01

Attention Driven Reference Resolution in Multimodal Contexts

John D. Kelleher

Technological University Dublin, john.d.kelleher@tudublin.ie

Follow this and additional works at: <https://arrow.tudublin.ie/scschcomart>

Part of the [Computer Engineering Commons](#)

Recommended Citation

Kelleher, J.: Attention Driven Reference Resolution in Multimodal Contexts. *Artificial Intelligence Review*. Vol. 25, issue 1-2, pp21-35. 2006. doi:10.1007/s10462-007-9022-9

This Article is brought to you for free and open access by the School of Computer Sciences at ARROW@TU Dublin. It has been accepted for inclusion in Articles by an authorized administrator of ARROW@TU Dublin. For more information, please contact arrow.admin@tudublin.ie, aisling.coyne@tudublin.ie, gerard.connolly@tudublin.ie.

Attention Driven Reference Resolution in Multimodal Contexts

J.D. Kelleher (john.kelleher@comp.dit.ie)*
School of Computing, Dublin Institute of Technology, Ireland

Abstract. In recent years a number of psycholinguistic experiments have pointed to the interaction between language and vision. In particular, the interaction between visual attention and linguistic reference. In parallel with this, several theories of discourse have attempted to provide an account of the relationship between types of referential expressions on the one hand and the degree of mental activation on the other. Building on both of these traditions, this paper describes an attention based approach to visually situated reference resolution. The framework uses the relationship between referential form and preferred mode of interpretation as a basis for a weighted integration of linguistic and visual attention scores for each entity in the multimodal context. The resulting integrated attention scores are then used to rank the candidate referents during the resolution process, with the candidate scoring the highest selected as the referent. One advantage of this approach is that the resolution process occurs within the full multimodal context, in so far as the referent is selected from a full list of the objects in the multimodal context. As a result situations where the intended target of the reference is erroneously excluded, due to an individual assumption within the resolution process, are avoided. Moreover, the system can recognise situations where attention cues from different modalities make a reference potentially ambiguous.

1. Introduction

Many modern natural language processing applications (human-robot collaboration, computer games, navigation aids etc.) share a visualised space with the user. In these applications the user interacts with the system using *situated language*. Situated language is spoken from a particular point of view within a physical or simulated context. The framework presented in this paper addresses a particular aspect of situated dialog, namely reference resolution.

A referring expression is a natural language expression that denotes an entity called a *referent*. Referring expressions come in a variety of *forms* including: definite descriptions, indefinites, pronouns, demonstratives. Each referring expression introduces a representation into the semantics of its utterance and this representation must be bound to an element in the context for the utterance's semantics to be fully resolved. Consequently, from a computational perspective reference resolution involves two main tasks:

* The author would like to thank the anonymous reviewers for their valuable comments and feedback.

1. creating and maintaining a model of the discourse context (this model should contain representations for all the entities that are available for reference)
2. matching/binding the representation introduced by a given referring expression to an element (or elements) in the set of possible referents

Most forms of referring expression have a preferred *mode of interpretation*: *anaphoric*, *exophoric*, etc. In a dialog, human participants expect their partner to construct and maintain a model of the evolving linguistic context. Referring expressions that access a representation in the linguistic context are interpreted anaphorically. However, in a situated dialog human participants expect their partner to not only construct and maintain a model of the linguistic discourse, but also to have full perceptual knowledge of the environment. Referring expressions that access a representation of an object that has not previously been referred to in the dialog but has entered the context through a non-linguistic modality (such as vision) are interpreted exophorically.

The following dialog excerpt, taken from the TRAINS-93 corpus (Allen and Schubert, 1991), illustrates the distinction between anaphoric and exophoric references. The excerpt is taken from a collaborative dialog between two participants, S1 and S2, who are trying to ship goods within a railroad freight system. Figure 1 illustrates the schematic representation of the railroad freight system that provided the visual context for the dialog. In this example, the indices i , j , k and l indicate that all the referring expressions marked by a particular index refer to the same entity.

1. **Visual context:** See Figure 1.

S1.1 “aha ... I see *an engine_i* and *a boxcar_j* both at *Elmira_k*”

S2.1 “right”

S1.2 “this looks like the best thing to do ... so we should get *the engine_i* to pick up *the boxcar_j* and head for *Corning_l* ... that sound reasonable”

S2.2 “sure ... that sounds good”

The references *an engine*, *a boxcar*, and *Elmira* in S1.1 and *Corning* in S1.2 are examples of exophoric references. The entities these expressions denote have not been previously mentioned in the dialog. As a result, these reference must be resolved relative to a set of entity representations in the context model that entered the model via the non-linguistic modalities, in this instance the visual context of the dialog. By contrast, the references *the engine* and *the boxcar* in S1.2 are examples of anaphoric references. The reference *the engine* can be resolved relative to the linguistic context by binding it to the representation of *an engine* introduced to the linguistic context

Figure 1. Map of TRAINS Domain

by the resolution of S1.1. Similarly, the reference *the boxcar* can be resolved relative to the linguistic context by binding it to the representation of *a boxcar* introduced by the resolution of S1.1.

However, there is no one-to-one relationship between form and mode of interpretation. For example, definite descriptions can be used either anaphorically or exophorically. Indeed, the two most common cases of definite descriptions in the TRAINS corpus of situated dialogue were anaphoric and exophoric definites (Poesio, 1993). One consequence of the one-to-many relationship between referential form and mode of interpretation, is that a multimodal reference resolution process should define a strategy to deal with cases where different mode of interpretations are suggested for the same reference. One solution, to this issue, is to define a preference ordering over the different interpretation rules. A second alternative is to use a probabilistic approach, where each interpretation of a reference is assigned a probability score that is used to rank the interpretations. The approach to reference resolution developed in this paper adopts a probabilistic solution to this issue. The framework uses the relationship between referential form and preferred mode of interpretation as a basis for a weighted integration of linguistic and visual attention scores for each entity in the context model. The resulting integrated attention scores are then used to rank the candidate referents during the resolution process, with the candidate scoring the highest selected as the referent.

One advantage of this approach is that resolution process occurs within the full multimodal context of the dialog, in so far as the referent is selected from a full list of the objects in the multimodal context ordered by a model of integrated salience. Consequently, none of the objects in the context are excluded from consideration. As a result situations where the intended target of the reference is erroneously excluded, due to an individual assumption within the resolution process, are avoided. Also, the framework can recognise cross-modal ambiguity by comparing the integrated salience of the primary candidate with the integrated salience of all the other objects in the context. In these ambiguous cases the initiation of a clarification dialog may be a

better system response rather than the selection of the primary candidate referent. By contrast, many of the previous multimodal resolution frameworks exclude entities in the multimodal context model from consideration before the selection of the referent. In some cases, for example Kievit et al. (2001), Salmon-Alt and Romary (2001), Landragin and Romary (2003), Kelleher et al. (2005), the initial set of candidate referents is restricted to a sub-set of the context based on preferences with respect to the mode of interpretation relative to the form of reference. In other frameworks, for example Gorniak and Roy (2004), candidate referents are incrementally excluded from consideration as the resolution process progresses due to the sequential manner that the semantics of the terms within the reference are processed.

Moreover, from a functional perspective this approach has the advantage of modularity and the potential to accommodate learning within the system. The modularity of the framework stems from the fact that the only information required by the resolution process from each of the information sources (language and vision) within the context are the attention scores for each entity. As a result, the resolution process is, to a large extent, decoupled from the representations and processes used within the linguistic and visual context models. The learning aspect of the system arises from the ease (relative to rule based approaches) with which the integration weightings associated with a particular form of reference could be updated, for example, using machine learning techniques such as reinforcement learning. Finally, from a cognitive perspective, an attention based model fits the theoretical and psychological data that points to the role of attention within human reference resolution. Grosz (1977) was first in observing the relationship between focus of attention and exophoric definite descriptions. More recently, psycholinguistic studies, such as Duwe and Strohner (1997), have shown that people often use perceptual salience to resolve linguistic references.

The paper is structured as follows: §2 reviews related work; §3 presents the data structures and algorithms used in the framework; §4 contains a worked example illustrating the functioning of the framework; the paper finishes, in §5, with conclusions.

2. Related Work

Grosz (1977) is arguably the seminal work on language and vision integration. This work highlighted that attention constrained and structured the processing of discourse. Moreover, Grosz was the first to observe the relationship between focus of attention and the use of exophoric definite descriptions: when an object is in the current mutual focus of attention it can be referred to by means of a definite description even though other objects fulfilling the

description have been introduced into the linguistic discourse or are present in the shared visible context.

Building on this work, Grosz and Sidner (1986) developed a focus stack model of global discourse attentional state. According to this model the common ground¹ can be divided into three parts: *the linguistic structure*, which contains information about the linguistic structure of utterances in the dialog; *the intentional structure*, which contains information about the goals of the participants in the conversation; and *the attentional structure*, which contains information about the objects introduced into the discourse and their relative salience. Furthermore, due to attentional constraints, discourse is segmented or chunked and when a definite description is used anaphorically, the only *antecedents*² considered are those in the same discourse segment.

Assuming Grosz and Sidner's (1986) focus stack model to be generally correct as a model of global discourse structure,³ the issue of how focus of attention and reference interact within a discourse segment must still be addressed. Several frameworks have been proposed, for example Alshawi (1987), Hajicová (1993), Lappin and Leass (1994) and Grosz et al. (1995).⁴ However, none of these models explicitly accommodate multimodal contexts.

Poesio (1993) reformulates the attentional model in Grosz and Sidner (1986) in situation theoretic terms. Interestingly, Poesio's framework separates the attentional common ground into several *anchoring resource situations*. For example, one anchoring resource is called the *discourse situation* and consists of a record of what has been said. This anchoring resource is used to interpret anaphoric references. Another anchoring resource situation called the *situation of attention* models the subset of information in the visual field of the discourse participants that they are attending to and is used to interpret exophoric⁵ definite descriptions. Furthermore, he defines rules within a default logic, called *principles for anchoring resource situations*, that predict whether a definite description is going to be interpreted anaphorically or exophorically. However, one of the issues with this approach is how to deal with conflicting defaults. Consequently, the framework cannot handle situations in which two principles of anchoring resource situations apply, one suggesting an anaphoric interpretation the other an exophoric interpretation.

Many computational frameworks for multimodal reference resolution have also been developed. McKevitt (1996) provides an excellent collection of papers on early systems. Recent systems that focus on multimodal reference resolution include: Kievit et al. (2001), Salmon-Alt and Romary (2001), Landragin and Romary (2003), Gorniak and Roy (2004) and Kelleher et al. (2005).

Kievit et al. (2001) define separate resolution strategies for each form of referring expression. A strategy consists of one or more resolution steps applied in a predefined order. A resolution step consists of 4 stages: (1) the selection of possible referents from a single sub-context (dialog, visual do-

main, etc.), (2) the filtering of this set of candidates, (3) the ordering of the candidates based on saliency, (4) an evaluation of the result. The algorithm halts as soon as one of the resolution steps finds a unique object or finds several objects and cannot choose which is the intended one. This approach is equivalent to a preference ordering being defined over the different modes of interpretation for each form of reference. One issue with this approach is that the set of candidates considered during any one resolution step is constrained to the set of entities within the sub-context the resolution step uses to construct the initial set of candidates. As a result, the system cannot recognise situations where a reference may be ambiguous between two entities in different sub-contexts, and, consequently, it may resolve a reference incorrectly rather than initiate a clarification process.

Gorniak and Roy (2004) focus on the resolution of references containing spatial descriptions. They propose a feed-forward filtering process to reference resolution. In their framework, each lexical item in the system's lexicon is associated with one or more *composer functions*. A composer function takes one or more candidate referents as input and filters this set of candidates by computing how well each of the candidates fulfills the semantic model defined for the lexical term. Reference resolution is carried out by chaining the composer functions associated with the lexical terms in the reference together, i.e. the filtered set of candidates output by one composer function is used as the input set by the next composer function in the chain. Gorniak and Roy note that this strategy can fail if one of the composer functions excludes the target object from the set of candidates. For example, when interpreting "the leftmost one in the front" the composer for "leftmost" selects the leftmost objects in the scene, not including the obvious example of "font" that is not a good example of "leftmost".

The reference resolution frameworks presented in Salmon-Alt and Romary (2001), Landragin and Romary (2003) and Kelleher et al. (2005) use the notion of a reference domain. A reference domain is a structured contextual subset of the multimodal dialog context. Reference domains are created in the context model due to perceptual or linguistic events or conceptual knowledge and are intended to reflect the mental representation of the event they model. In these frameworks the resolution process involves: (1) the construction of an underspecified reference domain, using templates associated with the form of the reference given; (2) the unification of this underspecified domain with a suitable reference domain within the context model; (3) the selection of one of the elements within the unified reference domain to function as the referent. However, similar to the frameworks proposed in Kievit et al. (2001) and Gorniak and Roy (2004), there is the potential for these frameworks to overcommit to a particular subset of the context during the resolution process. As the resolution process occurs within a sub-context, whose selection is at least partially driven by the form of the reference being interpreted, if the

wrong reference domain is selected the intended target object and/or plausible distractor referents, that may indicate the need for reference clarification, may be excluded from consideration.

3. Approach

Several theories of discourse reference have attempted to provide an account of the relationship between types of referential expressions on the one hand, and degrees of mental activation of discourse referents on the other (e.g. Ariel (1990), Gundel et al. (1993), Grosz et al. (1995)). A common theme among these accounts is that referential expressions need more coding material as the referent is less activated. Following these theories, the basic approach of the framework is to treat a given referring expression as a set of instructions that specifies how the spread of attention across the set of objects within the discourse context should be modified before the selection of the referent. Consequently, the concept of attention is at the core of the framework.

Studies of attention, for example Enns and Rensink (1990), Spivey-Knowlton et al. (1998), Hopfinger et al. (2000), Chum and Wolfe (2001), indicate that both bottom-up and top-down processes affect it. Bottom-up processing guides attention based on low-level perceptual cues. Top-down processing, driven by factors such as intention, also affect attention. Indeed, the results of several eye-tracking experiments, for example Yarus (1967), Spivey-Knowlton et al. (1998), Tanenhaus et al. (1995), indicate that language comprehension is one of the top-down processes affecting visual attention.

A concept closely related to attention is salience. In this paper salience is used to describe the factors and associated processes that direct attention. The framework distinguishes between three levels of salience.

Level 1 This level includes the basic visual salience (i.e. the prominence of an object due to bottom up visual cues) and linguistic salience (i.e. the prominence of an object due to previous discourse) of an object.

Level 2 This level consists of reference relative visual and linguistic salience. These saliency scores represent the salience of an object within each of the modalities within the context provided by the referring expression that is being resolved.

Level 3 This level represents the integrated salience of objects within the context provided by the referring expression that is being resolved. This is computed using a weighted combination of the object's level 2 salience scores. The weights used in this combination reflect the biasing associated with different forms of reference toward a particular information source.

1. compute the reference relative saliences for each object in the context
2. compute the integrated salience for each object in the context
3. return the object with the highest overall salience as the referent

Algorithm 1. Reference Resolution Algorithm

The flow of information during reference resolution is from level 1 to level 3. Algorithm 1 lists the basic steps in reference resolution.

In the following sections we describe the data structures, basic visual and linguistic salience algorithms used by the framework, and the algorithms used during reference resolution.

3.1. DATA STRUCTURES

The basic data structure used by the framework is called an **coreference class**. Each coreference class stores the saliency information for one object in the context model. Figure 2 illustrates the internal structure of a coreference class. The coreference class `id` is a unique string identifier. Each coreference class contains components for storing the basic and reference relative visual and linguistic salience scores and the integrated salience scores for the object the class represents in the context model.

New coreference classes are added to the context model as a result of visual processing. Each time an object is detected in the visual scene the context model is queried for the coreference class representing the object. If there is no coreference class for the object model a new coreference class is created and is assigned the id used by the vision processing. The basic visual salience component is initialised to the value created by vision processing when the object was detected. This is updated after each scene is rendered. All the other salience scores are initialised to 0. These components are updated after each utterance has been processed. Coreference classes are removed from the context model when both their basic visual and linguistic saliences fall below a threshold (.0001). In the following sections the algorithms that provide and use the information stored in these structures are described.

3.2. MODELLING BASIC VISUAL SALIENCE

Most computational models of visual attention focus on bottom-up processing, see Koch and Itti (2001) and Heinke and Humphreys (2004) for recent reviews. In most of these models several feature maps (such as colour, intensity etc.) are computed in parallel across the visual field and these are then combined into a single saliency map. Then a selection process deploys attention to locations in decreasing order of salience. In Kelleher and van Genabith (2004) a simple model of visual salience (based on object size and

id = String value
visual salience
a) basic = $[0 \dots 1]$
b) reference-relative = $[0 \dots 1]$
linguistic salience
a) basic = $[0 \dots 1]$
b) reference-relative = $[0 \dots 1]$
integrated salience = $[0 \dots 1]$

Figure 2. A Coreference Class

centrality relative to a focus of visual attention) was presented. In this paper we adopt use this model to capture the information entering the discourse through vision.

The visual salience algorithm uses a **false-colouring** technique. Each object in the simulation is assigned a unique colour or **vision-id**. This colour differs from the normal colours used to render the object in the world; hence the term false colouring. Each frame is rendered twice: firstly, using the objects' normal colours, textures and shading, and secondly, using the vision-ids. The first rendering is on screen (i.e. the user sees it), the second rendering may be off screen. After each frame is rendered, a bitmap image of the false colour rendering is created. The bitmap is then scanned and a list of the colours in the image is created. Using this list the system can recognise which objects are visible and which are not. Moreover, the system can identify, at the pixel level, the area covered by each object in the scene. This pixel information is used to compute the basic visual saliency of each object.

Mimicking the spread of visual acuity across the retina, the algorithm weights each pixel in the image based on its distance from the point of visual focus. The weighting is computed using Equation 1. In this equation, D equals the distance between the pixel being weighted and the point of focus, M equals the maximum distance between the point of focus and any point on the border of the image. The point of focus can be determined using eye tracking technology to compute the user's gaze at each scene rendering. However, if eye tracking is not being used the point of focus defaults to the center of the image or to the center of silhouette of the last object referred to. Algorithm 2 lists the procedure used to compute basic visual saliency and to update the coreference classes. For each scene processed the algorithm returns a list of objects in the scene each with a relative saliency between 0 and 1, with 1 representing maximum saliency.

```

for each object  $O_i$  in the scene do
   $AW(O_i)$  = average weighting of the pixels covered by  $O_i$ 
   $Total_{AW} = Total_{AW} + AW(O_i)$ 
endfor
for each coreference class  $CR_i$  do
  if  $CR_i$  is the coreference class representing  $O_i$  then
     $CR_i.basic\_visual = (CR_i.basic\_visual/2) + (AW(O_i)/Total_{AW})$ 
  else
     $CR_i.basic\_visual = CR_i.basic\_visual/2$ 
  endif
   $Total_{bvs} = Total_{bvs} + CR_i.basic\_visual\_salience$ 
endfor
for each coreference class  $CR_i$ 
   $CR_i.basic\_visual = CR_i.basic\_visual/Total_{bvs}$ 
endfor

```

Algorithm 2. The basic visual salience algorithm.

$$Weighting = 1 - \left(\frac{D}{M + 1} \right) \quad (1)$$

3.3. MODELLING BASIC LINGUISTIC SALIENCE

The basic linguistic salience of objects in the context are computed using an algorithm that is similar to Krahmer and Theune (2002). The algorithm is based on the ranking of the so-called **forward looking centers** (C_f) of an utterance. The set of forward looking centers of an utterance contains the objects referred to in that utterance. This set is partially ordered to reflect the relative prominence of the referring expressions within the utterance. Grammatical roles are a major factor here, so that `subject` > `object` > `other`. The central component of the algorithm is a function sf that maps the objects in a domain D to the set $\{0, \dots, 1\}$, with the intuition that 0 represents non-salience and 1 maximal salience. Figure 3 defines the salience function sf used by the framework. The algorithm assumes that in the initial situation s_0 all the objects in the domain are equally (not) salient: $sf(s_0, d) = 0$ for all $d \in D$.

It is not claimed that the function sf is the best way to assign linguistic salience. However, it does provides a reasonable, transparent and operational model of linguistic salience. Algorithm 3 defines the procedure used to update linguistic salience after an utterance has been processed.

Let U_i be a sentence uttered in state s_i , in which reference is made to $\{d_i, \dots, d_n\} \subseteq D$. Let $C_f(U_i)$ (the forward looking center of U_i) be a partial order defined over $\{d_i, \dots, d_n\} \subseteq D$. Then the salience weight of objects in s_{i+1} is determined as follows:

$$sf(s_{i+1}, d) = \begin{cases} 1 & \text{if } d = \text{subject}(U_i) \\ (sf(s_i, d)/2) + .5 & \text{if } d = \text{object}(U_i) \\ (sf(s_i, d)/2) + .25 & \text{if } d = \text{other}(U_i) \\ sf(s_i, d)/2 & \text{if } d \notin C_f(U_i) \end{cases}$$

Figure 3. Linguistic Salience Weight Assignment

```

Let  $Total_{DS} = 0$ 
for each coreference class  $CR_i$  do
     $CR_i.BLS = sf(s_j, CR_i)$ 
     $Total_{DS} = Total_{DS} + CR_i.BLS$ 
endfor
for each coreference class  $CC_i$  do
     $CR_i.BLS = CR_i.BLS / Total_{DS}$ 
endfor

```

Algorithm 3. The basic linguistic salience algorithm. BLS = basic linguistic salience.

3.4. COMPUTING REFERENCE RELATIVE SALIENCES

The first step in resolving a reference is to compute for each object in the context the salience of that object within each modality within the context provided by the reference. These *reference relative saliences* are computed for each object by integrating each object's basic visual and linguistic saliences with a rating of how well the object fulfills the selectional preferences⁶ encoded in the reference.

The rating of how well an object fits the description provided by a reference is called an *f-score*. Two f-scores are computed for each object for each reference: a visual and a linguistic f-score. Currently, the system can rate objects relative to their type, colour, size⁷ and location.⁸ Table I lists the ratings ascribed to an object for each type of selectional preference. An object's visual f-score is initialised to 0 and its ratings are integrated using addition. An object's linguistic f-score is initialised to 1 and its ratings are integrated using multiplication.

Once the f-scores have been computed the object's reference relative visual and linguistic saliences are computed by integrating the f-scores with its basic visual and linguistic salience. Again, addition is used for integration in the visual context and multiplication is used for integration in the linguistic context. Consequently, an object's reference relative visual salience will be

Table I. Selectional Preferences Scores

	Fulfills	Not Fulfill
TYPE	1	0
COLOUR	1	0
SIZE	[1 ... 0]	
LOCATION	[1 ... 0]	

```

for each coreference class  $CR_i$  do
   $f\_score_{linguistic} = 1$ 
   $f\_score_{vision} = 0$ 
  for each selectional preference  $sp_j$  in the description do
     $f\_score_{linguistic} = f\_score_{linguistic} * rating(CR_i, sp_j)$ 
     $f\_score_{vision} = f\_score_{vision} + rating(CR_i, sp_j)$ 
  endfor
   $CR_i.RLS = CR_i.BLS * f\_score_{linguistic}$ 
   $Total_{rls} = Total_{rls} + CR_i.RLS$ 
   $CR_i.RVS = CR_i.BLS * f\_score_{vision}$ 
   $Total_{rvs} = Total_{rvs} + CR_i.RVS$ 
endfor
for each coreference class  $CR_i$  do
   $CR_i.RLS = CC_i.RLS / Total_{rls}$ 
  STATE  $CC_i.RVS = CC_i.RVS / Total_{rvs}$ 
endfor

```

Algorithm 4. Computing the reference relative saliences. RLS = reference relative linguistic salience, RVS = reference relative visual salience, BLS = basic linguistic salience, BVS = basic visual salience.

> 0 if it fulfils any of the selectional preferences in the description, and its linguistic reference relative salience will be = 0 if it does not fulfill all of the selectional preferences in the description. Algorithm 4 lists the algorithm for computing the reference relative saliences.

3.5. CREATING THE INTEGRATED CONTEXT AND SELECTING THE REFERENT

The final step before the selection of the referent is the integration of each object's reference relative saliences. This is done using a weighted combination. The weightings are dependent on the form of referring expression (e.g. definite descriptions versus pronominal references) being resolved and reflect the preferential interpretation associated with each type of reference. For ex-

```

for each coreference class  $CR_i$  in the context model do
  Let  $index = 0$ ,  $max = 0$ ,  $interval = 0.3$ 
  if reference = definite description then
     $CR_i.integrated = (CR_i.RVS * 0.5) + (CR_i.RLS * 0.5)$ 
  elseif reference = pronominal reference then
     $CR_i.integrated = (CR_i.RVS * 0.1) + (CR_i.RLS * 0.9)$ 
  endif
  if  $CR_i.integrated > max$  then
     $index = i$ 
     $max = CR_i.integrated$ 
  endif
endfor
for each coreference class  $CR_j$  in the context model do
  if  $j \neq index$  then
    if  $CR_j.integrated > CR_{index}.integrated - interval$  then
      return 0 //Reference Deemed Ambiguous
    endif
  endif
endfor
return  $CR_{index}$ 

```

Algorithm 5. Constructing the integrated context and selecting the references. RVS = reference relative visual salience, RLS = reference relative linguistic salience.

ample, in general, a pronoun is used to refer to a referent that is prominent within the linguistic context. By contrast, a definite description can be used to refer to an object from the visual scene and to previously mentioned objects. Currently, the system uses predefined weights for this integration. When resolving a definite description visual and linguistic salience are integrated evenly. When resolving a pronominal reference the integration weightings used biases towards linguistic salience. Algorithm 5 defines the procedure used to construct the integrated context and select the reference. It also defines the mechanism used to check for ambiguous references. This ambiguity check uses a predefined confidence interval and simply checks that within the context provided by the referring expression the integrated salience of the object selected as the referent is sufficiently larger than the other objects in the context to ensure that the reference is not ambiguous. In situations where the ambiguity check fails the algorithm returns 0.

4. Worked Example

The functioning of the framework can be illustrated using a worked example. The example uses Figure 4 as a visual context, and the utterances (1) and (2) as the example discourse.

Figure 4. Example visual context. H1 = red house, H2 = green house.

1. make *the red house* green
2. make *the tree to the left of it* bigger

Table II lists the saliences scores computed by the framework during the different stages of this interaction. Rows 1 and 2 of the table present the initial basic visual and linguistic salience scores of the objects in Figure 4. Rows 3 to 7 presents the f-scores and reference and integrated saliences computed for the objects when the system processed *the red house*. The asterisk in line 7, H1's column, indicates the highest integrated salience at the end of the resolution process. As a result of obtaining the maximum salience H1 is selected as the referent. Rows 8 and 9 list the basic salience scores for the objects after the basic linguistic salience has been updated and the point of visual focus has been located at the center of H1's silhouette. The movement of the visual focus away from the center of the image is reflected in the increases in the basic visual salience of H1 ($0.3271 \rightarrow 0.3581$) and T1 ($0.1728 \rightarrow 0.2938$). Rows 10 to 14 list the f-scores and reference and integrated saliences computed for the objects when the system processed *it*. The biasing towards linguistic salience is apparent in the dominance of H1's integrated salience. Rows 15 to 19 list the f-scores and reference and integrated saliences computed for the objects when the systems processed *the tree to the left of it*. The difference between the visual and linguistic f-scores of T1 and T2 is due to the locational description: T1 was judged by the system to fulfill the locational description with a rating of 0.9396, while T2 was judged not to fulfill the description and was ascribed a rating of 0.0000 for this selectional preference. As a result, T1 achieved the highest salience (0.5669) and was selected as the referent. Figure 4 illustrates the visual context at the end of the interaction.

5. Conclusions and Future Work

This paper presented an attention based reference resolution framework for visually situated discourse. The framework uses a weighted integration of

Table II. Saliency scores computed during the example interaction.

		H1	H2	T1	T2
Initial Context					
1	Basic visual salience	0.3271	0.3272	0.1728	0.1728
2	Basic linguistic salience	0.0000	0.0000	0.0000	0.0000
the red house					
3	Visual f-score	2.0000	1.0000	0.0000	0.0000
4	Linguistic f-score	1.0000	0.0000	0.0000	0.0000
5	Reference visual salience	0.5818	0.3318	0.0432	0.0432
6	Reference linguistic salience	0.0000	0.0000	0.0000	0.0000
7	Integrated salience	0.5818*	0.3318	0.0432	0.0432
8	Basic visual salience	0.3581	0.2273	0.2938	0.1208
9	Basic linguistic salience	1.0000	0.0000	0.0000	0.0000
it					
10	Visual f-score	0.0000	0.0000	0.0000	0.0000
11	Linguistic f-score	1.0000	1.0000	1.0000	1.0000
12	Reference visual salience	0.3581	0.2273	0.2938	0.1208
13	Reference linguistic salience	1.0000	0.0000	0.0000	0.0000
14	Integrated salience	0.9358*	0.0227	0.0293	0.0120
the tree to the left of it					
15	Visual f-score	0.0000	0.0000	1.9396	1
16	Linguistic f-score	0.0000	0.0000	0.9396	0
17	Reference visual salience	0.0909	0.0577	0.5669	0.2845
18	Reference linguistic salience	0.0000	0.0000	0.0000	0.0000
19	Integrated salience	0.0909	0.0577	0.5669*	0.2845

visual and linguistic attention to order the candidate referents within the context. The candidate with the highest integrated attention score is taken to be the referent. One advantage of this approach is that the resolution process occurs within the full multimodal context. As a result situations where the intended target of the reference is erroneously excluded, due to an individual assumption within the resolution process, are avoided. Moreover, the system can recognise situations where attentional cues from different modalities make a reference potentially ambiguous. From a cognitive perspective the

Figure 5. The final state of the visual context.

framework meshes well with psycholinguistic results that point to the role of attention within human reference resolution processes.

Finally, it should be noted that the framework as it currently stands is intended to represent an abstract and preliminary attempt. Several issues need to be addressed if it is to be used as a component within a dialog systems for less constrained contexts. In particular, the use of predefined weights for salience integration is overly simplistic. This issue could be addressed by using a machine learning algorithm, such as reinforcement learning, to automatically compute these weights. The visual and linguistic salience algorithms should also be improved. For dialog systems interfacing with virtual environments, the visual salience algorithm should be extended to at least handle attentional cues such as color, motion and location of gaze. If the framework was to be used within a real-world system, such as a robot dialog system, a computer vision saliency algorithm, such as Itti and Koch (2000), could be adopted. The linguistic saliency algorithm should also be revised and extended. For example, an information based approach to linguistic attention, such as Hajicová (1993), may be more suitable than the simplified Centering (Grosz et al., 1995) based framework proposed here. Moreover, the relationship between the framework's model of local level attention and a more global model of discourse structure, such as Grosz and Sidner's (1986) focus stack model or Asher and Lascarides' (2003) SDRT framework, should be clarified. Fortunately, the modular nature of the framework makes such modifications possible without major changes to the overall approach.

Notes

¹ The dialog participants mutually developed public view of what they are talking about.

² The antecedent of an anaphoric reference is the representation of the reference's referent that was introduced to the discourse model by a prior referring expression.

³ For alternate models see (Hobbs, 1985; Mann and Thompson, 1987; Asher and Lascarides, 2003).

⁴ See Kruijff-Korbyová and Hajicová (1997) for a comparison of these approaches.

- ⁵ Poesio uses the term *visible situation use* to describe to exophoric definite descriptions.
- ⁶ The semantics of the descriptive terms used in the reference.
- ⁷ An objects size rating is based on the number of pixels it covers relative to the other objects in the scene.
- ⁸ An objects location rating is computed using the AVS model described in Regier and Carlson (2001)

References

- Allen, J. and L. Schubert: 1991, 'The TRAINS Project'. Technical report, University of Rochester, Department of Computer Science.
- Alshawi, H.: 1987, *Memory and Context for Language Interpretation*. Cambridge, UK: Cambridge University Press.
- Ariel, M.: 1990, *Accessing noun-phrase antecedents*. London: Routledge.
- Asher, N. and A. Lascarides: 2003, *Logics of Conversation*. Cambridge University Press.
- Chum, M. and J. Wolfe: 2001, 'Visual Attention'. In: E. B. Goldstein (ed.): *Blackwell Handbook of Perception*, Handbooks of Experimental Psychology. Blackwell, Chapt. 9, pp. 272–310.
- Duwe, I. and H. Strohner: 1997, 'Towards a Cognitive Model of Linguistic Reference'. Report: 97/1 - Situierete Künstliche Kommunikatoren 97/1, Universität Bielefeld.
- Enns, J. and Rensink, R.A.: 1990, 'Influence of scene-based properties on visual search'. *Science* **247**, 721–723.
- Gorniak, P. and D. Roy: 2004, 'Grounded Semantic Composition for Visual Scenes'. *Journal of Artificial Intelligence Research* **21**, 429–470.
- Grosz, B.: 1977, 'The Representation and Use of Focus in Dialogue Understanding'. Ph.D. thesis, Stanford, University.
- Grosz, B., A. Joshi, and W. Weinstein: 1995, 'Centering: A Framework for Modelling Local Coherence of Discourse'. *Computational Linguistics* **21**(2), 203–255.
- Grosz, B. and C. Sidner: 1986, 'Attention, intentions, and the structure of discourse'. *Computational Linguistics* **12**(3), 175–204.
- Gundel, J., N. Hedberg, and R. Zacharski: 1993, 'Cognitive Status and the form of referring expression in discourse'. *Language* **69**, 274–307.
- Hajicová, E.: 1993, *Issues of Sentence Structure and Discourse Patterns*, Vol. 2 of *Theoretical and Computational Linguistics*. Charles University Press.
- Heinke, D. and G. Humphreys: 2004, 'Computational Models of Visual Selective Attention: A Review'. In: G. Houghton (ed.): *Connectionist Models in Psychology*. Psychology Press.
- Hobbs, J.: 1985, 'On the coherence and structure of discourse'. Technical Report CSLI-85-37, Center for the Study of Language and Information.
- Hopfinger, J., M. Buonocore, and G. Mangun: 2000, 'The neural mechanisms of top-down attentional control'. *Nature Neuroscience* **3**(3), 284–291.
- Itti, L. and C. Koch: 2000, 'A Saliency-Based Search Mechanism for Overt and Covert Shifts of Visual Attention'. *Vision Research* pp. 1489–1506.
- Kelleher, J., F. Costello, and J. van Genabith: 2005, 'Dynamically structuring, updating and interrelating representations of visual and linguistic discourse context'. *Artificial Intelligence* **167**(1-2), 62–102.
- Kelleher, J. and J. van Genabith: 2004, 'Visual Saliency and Reference Resolution in Simulated 3D Environments'. *AI Review* **21**(3-4), 253–267.
- Kievit, L., P. Piwek, R. Beun, and H. Bunt: 2001, 'Multimodal Cooperative Resolution of Referential Expressions in the DenK System'. In: H. Bunt and R. Beun (eds.): *Cooper-*

- ative Multimodal Communication*, Lecture Notes in Artificial Intelligence 2155. Berlin Heidelberg: Springer-Verlag, pp. 197–214.
- Koch, C. and L. Itti: 2001, ‘Computational Modelling of Visual Attention’. *Nature Reviews Neuroscience* **2**(3), 194–203.
- Krahmer, E. and M. Theune: 2002, ‘Efficient Context-Sensitive Generation of Referring Expressions’. In: K. van Deemter and R. Kibble (eds.): *Information Sharing: Reference and Presupposition in Language Generation and Interpretation*. Stanford: CLSI Publications.
- Kruijff-Korbayová, I. and E. Hajicová: 1997, ‘Topics and Centers: A Comparison of the Salience-Based Approach and the Centering Theory’. *Prague Bulletin of Mathematical Linguistics* **67**, 25–50. Charles University, Prague, Czech Republic.
- Landragin, F. and L. Romary: 2003, ‘Referring to Objects Through Sub-Contexts in Multimodal Human-Computer Interaction’. In: *DiaBruck 7th workshop on the semantics and pragmatics of dialogue, Sept 4th-6th 2003*. University of Saarland, Germany.
- Lappin, S. and H. Leass: 1994, ‘An Algorithm for Pronominal Anaphora Resolution’. *Computational Linguistics* **20**(4), 535–561.
- Mann, W. and S. Thompson: 1987, ‘Rhetorical structure theory: Description and construction of text structures’. In: G. Kempen (ed.): *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*. Dordrecht: Nijhoff., pp. 83–96.
- McKevitt, P. (ed.): 1995/1996, *Integration of Natural Language and Vision Processing (Vols. I-IV)*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Poesio, M.: 1993, ‘A situation-theoretic formalization of definite description interpretation in plan elaboration dialogues’. In: P. Aczel, D. Israel, Y. Katagiri, and S. Peters (eds.): *Situation Theory and its Applications*, Vol. 3. pp. 339–374, CSLI.
- Regier, T. and L. Carlson: 2001, ‘Grounding Spatial Language in Perception: An Empirical and Computational Investigation’. *Journal of Experimental Psychology: General* **130**(2), 273–298.
- Salmon-Alt, S. and L. Romary: 2001, ‘Reference resolution within the framework of cognitive grammar’. In: *Proceedings of the Seventh International Colloquium on Cognitive Science (ICCS-01)*. Donostia, Spain, pp. 284–299.
- Spivey-Knowlton, M., M. Tanenhaus, K. Eberhard, and J. Sedivy: 1998, ‘Integration of Visuospatial and Linguistic Information: Language Comprehension in Real Time and Real Space’. In: P. Olivier and K. Gapp (eds.): *Representation and Processing of Spatial Expressions*. Lawrence Erlbaum Associates, pp. 201–214.
- Tanenhaus, M., M. Spivey-Knowlton, K. Eberhard, and J. Spivey: 1995, ‘Integration of Visual and Linguistic Information in Spoken Language Comprehension’. *Science* **268**, 1632–1634.
- Yarbus, A.: 1967, *Eye Movements and Vision*. New York: Plenum Press.