

WHO ARE YOU?

On the Acquisition of Information about People for an Agent that Remembers

Nikita Mattar and Ipke Wachsmuth

Artificial Intelligence Group, Bielefeld University, Bielefeld, Germany

Keywords: Conversation topics, Embodied conversational agents, Human-agent interaction, Long-term interactions, Person memory, Small talk, Social categories, Stereotypes.

Abstract: Humans make extensive use of specialized representations to remember people they interacted with. While current research on embodied conversational agents focuses on the relationship between agent and interlocutor, the representation of the latter is mostly neglected. But information about others are inevitable for an agent to adapt to its interlocutors and to establish long-term relationships with them. In this work, we present a model of Person Memory for virtual agents. We discuss what kinds of information have to be stored about people. Furthermore, we stress the importance of social categories. In our scenario, we focus on first encounters between our agent and people. We show how the agent is able to exploit his Person Memory to acquire information about others during Small Talk and guide the conversation.

1 MOTIVATION

When two people meet for the first time they often would like to know from each other: Who are you? How do humans acquire information of others they encounter for the first time? How and what information do they remember from each other?

Research on human-like memory for virtual agents has gained a lot of attention recently. Agents equipped with autobiographical memory, e.g. (Kasap et al., 2009), or episodic memory, cf. (Brom et al., 2007), are enabled to remember their own experiences. Companion agents greatly benefit from these memories, in that information from previous interactions can be accessed and used by the agent. However, in the above approaches the memories are centered on the agent's experiences, yet people the agent interacted with play a minor role.

In this work we present a model of Person Memory for an embodied conversational agent that focuses on the representation of people. Our agent MAX has a cognitive architecture based on BDI. The agent resides on the hallway of the AI group of Bielefeld University. A broad range of Small Talk knowledge allows for short enjoyable interactions with him. But over the long run, conversations with our agent are more or less the same. So it is not easy to build up some kind of relationship over time, since people may get annoyed by repetitive interactions. To enable our agent to build up longer lasting relationships, he

should be put in the position to adapt to his interlocutors. In that, he must be able to distinguish between different kinds of people, in order to know what to talk about. To remember past interactions would allow him to avoid repetition of things already said, and to pick up interesting topics of previous conversations.

Small Talk is considered important to increase trust and familiarity (Bickmore and Cassell, 2001) between virtual agents and their interaction partners. Yet, most of the systems equipped with small-talk abilities are restricted to common topics, like the current weather. Small Talk about topics of interest to certain individuals has been mostly neglected. In this paper, we show how information of social categories can be exploited in order to acquire information about, and guide a conversation with, a new acquaintance.

The paper is organized as follows. In Section 2 related work is reviewed. In order to describe the Person Memory component, in Section 3, we review work of social psychology, dealing with person perception and social categories. The architecture and Person Memory is described in Section 4. In Section 5 we present our scenario, and show how social categories can be used to determine topics of interest to different kinds of people. A Summary and future work conclude this paper.

2 RELATED WORK

Stereotypes have been found an adequate way in recommender systems in order to tailor recommendations to different kinds of people. E.g. GRUNDY (Rich, 1979) demonstrated how stereotypes can be used to recommend books to people. GRUNDY prompted the user for some personality traits at the beginning of the interaction. It then used this information to assign stereotypes to be able to give recommendations on books.

Since GRUNDY, a lot of research in the area of recommender systems has been done. Goal of these systems is to provide recommendations on e.g. news, books, or movies that match the interest of their users as close as possible. Recommender systems are usually classified into content-based, collaborative, and hybrid approaches. For an overview refer to (Adomavicius and Tuzhilin, 2005).

Systems that provide user models for recommender systems developed into powerful tools. Kobsa provides an overview of generic user modeling systems (Kobsa, 2007). But, conversational agents differ from traditional recommender systems and their user models. Most conversational agents are designed to support users in certain tasks, like learning, diet planning etc., thus the relation between agent and interaction partner plays an important role.

Research on relational agents focuses on the establishment of long-term relationships with their users. In (Bickmore and Cassell, 2001) an extended version of REA, an embodied conversational agent, is described. Besides of task-related talk, REA is capable of engaging in Small Talk to deepen the relationship to her interlocutors. While building and maintaining long-term relationships, and even reidentifying users (Schulman et al., 2008), is a central part, little information is given on how persons the agents interact with are represented.

The relationship between agent and interlocutor has shown to be improved when the agent engages in Small Talk (Campbell et al., 2009). But, the complex structure of Small Talk has been mostly neglected. In a recent work of (Endrass et al., 2011) cultural differences in Small Talk are investigated and implemented for multi-agent systems. Following (Schneider, 1988), they categorize Small Talk topics into three different categories: immediate, external, and communication situation.

More recently, systems with human-like memory gained attention. Agents equipped with autobiographic or episodic memories are able to remember past events they experienced either in a virtual world on their own (Brom et al., 2007), or with people they

interacted with (Kasap et al., 2009). The agent described in the latter, e.g. stores the user and the change of relationship in its episodic memories. However, again there is no information on how the user is represented by the agent. In general, besides the fact that user modeling is commonly employed in recommender systems, most conversational agents lack a clear representation of the people they interact with.

3 PERSON PERCEPTION

In order to equip a virtual agent with a representation of people several questions need to be addressed. How do humans perceive others? Is there a difference between the representation of people and other world objects in humans? How do humans acquire knowledge about others they have not met before? What do humans remember about other people?

Humans try to simplify the surrounding world by making use of categories. Categorization helps to focus on important things by reducing information. Not only do humans categorize world objects like cups and cars, but we categorize other people we meet as well (Stangor and Lange, 1994).

While natural objects are mostly perceived through their form and function (Rosch as cited in Dahlgren, 1985), members of social groups are perceived in a more complex way. E.g., it is not always possible to conceive a social category by observing someone, but the social role of the individual needs to be clear. Furthermore, more information is stored in the representation of a social category, such as social function, relation between members of categories, and internal traits (Dahlgren, 1985).

Social categories consist of stereotypical information humans have about a certain group of people (Hamilton, 1979). Humans are able to assign some social categories even before actually talking to somebody, e.g. categories like age or gender. The impression of the other is refined when more information is acquired, e.g. during conversation. Information like occupation, place of residence, and individual interests allow for a fine grained categorization. Social categories are used as a basis of encoding and the associated information influence judgment of and behavior towards others (Stangor and Lange, 1994).

To sum up, in order to represent people in a cognitive architecture of a virtual agent, it seems not sufficient to only consider individuals the agent has interacted with. Person perception in humans is strongly influenced by knowledge one has about groups of people, i.e. social categories. To equip an agent with such knowledge will enhance its capabilities of react-

ing to different kinds of people the agent has not met before, and the storage of information about people, as well.

3.1 What to Remember about People

We identified the following five types of information to be stored about people which form the basis of our conception of person memory:

Biographical Facts. During conversation people reveal facts about themselves, like their age, occupation, hometown, etc. These kinds of information can be compared to autobiographical facts (Conway, 1987) one remembers about oneself.

Preferences and Interests. Preferences and interests are valuable when trying to find topics to talk about or that should be avoided during conversation. Furthermore, they can be used to determine similarity between people.

Personality Traits. Information about the personality of the interlocutor allows to adapt the agent's personality. It has been shown that people prefer to talk to an agent that matches their personality (Bickmore and Cassell, 2001).

Events and Episodes. When we think about a specific person previous events one experienced with the other come to mind. In order to access significant events the agent experienced with someone, links to these events should be stored in the representation of a person. This enables the agent to bring up a previous encounter that was significant for their relationship during conversation.

Relationship Information. Information about the relationship between the agent and a person enables the agent to determine appropriate topics for Small Talk (Endrass et al., 2011). Furthermore, the agent's emotions are especially influenced by the relationship to an interlocutor: the agent might be happier when meeting someone he likes than meeting someone he has never met before.

4 A PERSON MEMORY FOR A VIRTUAL AGENT

To address these topics, the cognitive architecture is extended by a Person Memory to allow virtual agents to keep track of people they interacted with, and an Event Memory (not discussed in this paper) for storage and retrieval of past interactions.

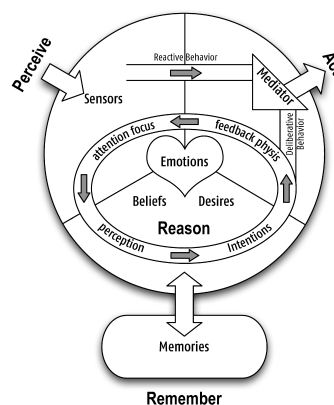


Figure 1: The cognitive architecture of the agent extended with a memory component.




Figure 2: Person and Event Memory of the cognitive architecture.

4.1 Cognitive Architecture

The cognitive architecture consists of a reactive and a deliberative perceive-reason-act part (figure 1 top). While the reactive part directly maps sensor information to behaviors, the deliberative part relies on a BDI-kernel based on JAM (Huber, 1999). The reactive and deliberative parts are executed concurrently and corresponding behaviors can overrule each other in order to control the agent (Lessmann et al., 2006). Working memory (WM) consists of JAM beliefs about the current situation and the world. In our architecture the term working memory has to be understood as a technical term. At present there is no limitation of information like in human working memory.

Since there was no clear distinction between WM and long-term memory (LTM) in the agent's architecture, we introduced a new memory component (figure 1 bottom). It allows access of long-term memories during reasoning and is composed of two different modules: Person Memory and Event Memory (figure 2). The memory modules are connected through a gatekeeper with the BDI architecture. The gatekeeper is able to select what information is passed from and to the working memory of the agent. PMPU and EMPU are the processing units of both modules

```

Agent: What is your name?
Visitor: My name is Paul.

<rule>
  <match>
    <convfunction type="provide.content.name" />
    <pattern>My name is $name</pattern>
  </match>
  <action>
    <command function="memorize"
      arguments="name:$name"/>
  </action>
</rule>

```

Figure 3: The agent asks his interlocutor for his name (top). Dialog rule matching the interlocutors response, with encoding and transferral of the name into working memory as corresponding action (bottom).




Figure 4: Types of access to WM and LTM and flow of information. a) The long-term memory can be queried for information by JAM rules. Available information is transferred into WM. b) Activation of related long-term memories occurs when information enters WM. The additional information is transferred into WM.

(Person Memory and Event Memory processing unit). They carry out retrieval, storage, and processing of information.

The dialog engine of our system follows a pattern matching approach. Utterances of the interlocutor are processed by JAM rules. If a rule matches the input the according JAM plan is executed. Figure 3 shows an example of the encoding and memory transferral process. A dialog rule matching the response of the interlocutor to the agents question triggers the transferral to WM. The information is encoded as a slot-value pair. Accordingly, there has to be a dialog rule for every piece of information our agent should be able to remember.

4.1.1 Activation

Incoming parts of information can lead to an activation of already memorized information (figure 4 b). In the conversation of figure 5 (top), a visitor provides his name to the agent. The information is transferred into working memory, passes the gatekeeper, and leads to an activation of information known about

```

Agent: Hello, I am MAX. What is your name?
Visitor: Hi MAX, It is me, Paul.
[Activation of additional LTM information and transferral to WM]
Agent: Paul from Bielefeld?
[Use of information from WM]
Visitor: Yes, that's right.
Agent: Nice to meet you again!

Visitor: How many people from Bielefeld do you know?
Agent: Hmm, I dont know.
[No information in WM]
Agent: Let me think about it!
[Query LTM for information]
Agent: I have met 6 people from Bielefeld.

```

Figure 5: Fragment of a conversation between our agent and a visitor the agent met before. Additional information is activated and transferred into WM (top). Information is explicitly queried from LTM (bottom).

a person with that name inside the long-term memory. The activated information is transferred into working memory. In the example, the agent is thus able to use the information and ask an additional question in order to clarify if the person activated in the Person Memory refers to the right one.

4.1.2 Recall

As not all the information stored in long-term memory is activated and transferred into WM at once, the agent needs a way to recall information form his long-term memories. In figure 5 (bottom), a short fragment of a conversation with our agent demonstrates a use-case of direct access to LTM. The agent does not have enough information in his WM to answer the question of the visitor. In this case he queries his LTM for the number of people he knows from Bielefeld and is able to answer the question (figure 4 a).

4.2 Representation of Social Categories

In the Person Memory, social categories are defined in a frame like manner. Table 1 depict excerpts of social categories of our system. Every slot represents a stereotypical information that is associated with the social category. Stereotypical information in social categories can consist of the first three types of information introduced in Section 3.1: Biographical facts, preferences and interests, and personality traits. In this work we focus on interests in certain topics.

A confidence value $c \in [0, 1]$ and a modifier value $m \in (-\infty, \infty)$ are associated with each stereotype. The confidence of a stereotypical information specifies

Table 1: Excerpts of social category frames: male, female, computer science student, and sports student.

Slot	Value	Conf.	Mod.
interestedin	soccer	0.6	2
interestedin	shopping	0.2	-5
...			

Slot	Value	Conf.	Mod.
interestedin	soccer	0.4	-3
interestedin	shopping	0.8	5
...			

Slot	Value	Conf.	Mod.
interestedin	c. games	0.6	5
interestedin	soccer	0.3	-4
interestedin	shopping	0.2	-3
...			

Slot	Value	Conf.	Mod.
interestedin	c. games	0.2	-5
interestedin	soccer	0.8	5
interestedin	shopping	0.7	3
...			

how probable it is that the information applies for a member of the category. This allows triggering certain behaviors if a stereotype proves wrong during conversation. E.g. the agent could be surprised if he talks to a computer science student who does not like to play computer games, but loves to play soccer, since his expectancies due to his social categories would be the opposite. The modifier allows altering the utility of a stereotypical information for conversation. How the utilities of information can be calculated are described in Section 5.1 .

4.3 Representation of Persons

The Person Memory stores information about the agent himself, people the agent met, and people the agent heard about. The representation of a person follows a frame-based approach, as well. Table 2 shows an empty person frame.

The initial person frame is sparse and only contains slots for a few important information about a person. The confidence value is used to annotate information acquired directly from the interlocutor during conversation (high confidence), and information that is e.g. inferred from other information (low confidence). The modifier, again, is used to alter the utility for an information for conversation.

Following (Rich, 1979), every social category is annotated with one or more facts that trigger their activation. If a fact about a person is revealed during

Table 2: Empty person frame.

Slot	Value	Conf.	Mod.
id			
firstname			
lastname			
gender			
age			
hometown			
social category			
...			
interest			

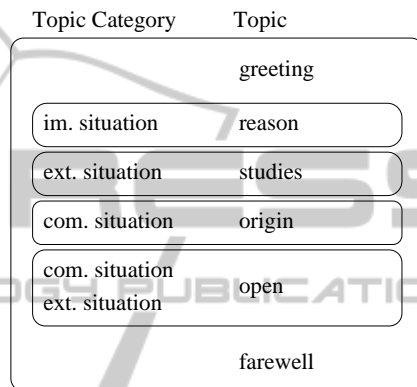


Figure 6: Structure of the testbed dialog. Topics are categorized in three categories: immediate, external, and communication situation.

conversation and triggers a social category, the category is assigned to the person. From this point, the stereotypes can be used by the agent, for example, to select topics that might be interesting to the interlocutor. Depending on the response of the interlocutor the stereotype might be confirmed or rejected and can be included in the corresponding person frame.

As one acquires specific knowledge about an individual, one is less likely to make use of the social categories associated with the person (Senay and Keysar, 2009). Therefore, further uses of this information will refer to the information in the person frame instead of the one in the social category.

5 SCENARIO

As a testbed for the Person Memory, we designed a dialog scenario. Since our agent resides on the hallway of our group, mostly students approach the agent during the day. In our scenario, our agent and a student meet on our hallway and have a short conversation, while waiting. To model our dialog, we analyzed the initial-meeting dialogs of participants of the CUBE-G corpus (Rehm et al., 2007), and used their structure as

a guideline.

Figure 6 depicts the structure of our dialog. Classification of topics follows the classification used in (Endrass et al., 2011). After the initial greeting, the agent asks the interlocutor why he is waiting, and then proceeds talking about the student's subject of study. Then the interlocutor is asked for his origin.

Up to this point the dialog topics are fixed. This enables the agent to acquire some initial social categories. Next, further topics can be discussed during an open slot in the dialog structure.

We now present how social categories can be exploited to determine topics that are suitable to talk about as the conversation continues.

5.1 Exploiting Social Categories

Social categories and the associated stereotypical information are well suited for Small Talk. Exploiting social categories, an ECA is able to ascribe attributes to people with very little information about them. This enables the agent to identify topics that might be interesting to someone, rather than suggesting what to talk about out of the blue.

To select suitable topics for conversation, not only the interests inferred from the social categories of the interlocutor have to be taken into account. Since we aim at building a human-like interaction partner, our agent is equipped with own interests that might interfere with the ones of his interaction partner. Therefore, the interests of the agent need to be considered when selecting a topic for conversation.

5.1.1 Combining Social Categories

For topic selection, stereotypical information of all social categories of a person are combined to find a suitable topic.

The utility $u(I)$ of an information I of a social category C is calculated according to (1).

$$u(I) = m(I) * c(I) \quad (1)$$

In equation (1), $m(I)$ denotes the modifier, $c(I)$ the confidence of an information I . The total utility value is acquired by adding all the single utility values of an information:

```
for (Information i : informations) {
    for (SocialCategory sc : categories) {
        utility_i += u( sc.get(i) );
    }
}
```

Figure 7 demonstrates how the utility values for each interest of a person change during conversation.

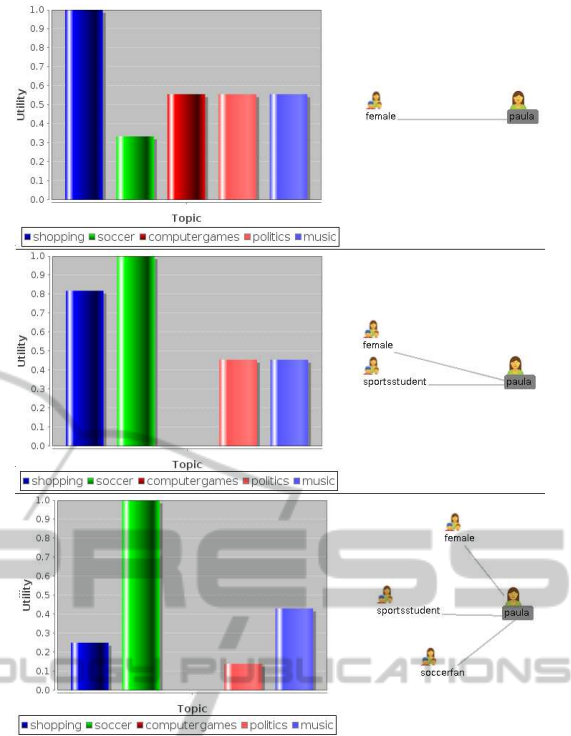


Figure 7: Adaptation of normalized utility values of interests of a person, as more social categories are being acquired.

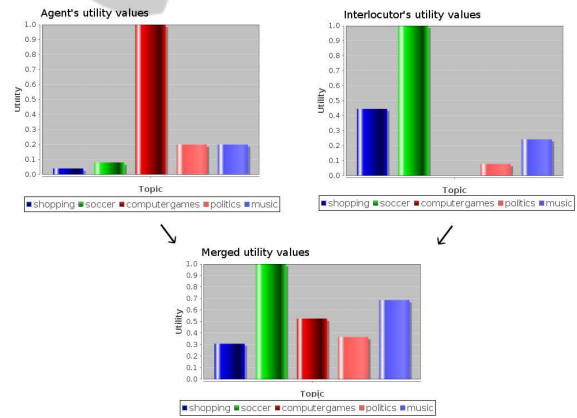


Figure 8: Utility values of the agent's, and an interlocutor's interests, and calculated final utility values.

5.1.2 Topic Selection

As stated above, not only the interlocutor's interests are considered, but the agent's interests as well. Furthermore, the course of conversation influences the utility of a topic. In that, a topic that was already talked about has a lower utility. The final utility $u_f(I)$ is calculated according to equation (2).

keep things going, has not been exploited to its full extent. Research on Small Talk shows that it is not as random as often perceived. It does follow certain rules, depends on the situational context, the relationship between the interlocutors, and goes further than chatting about the weather. We demonstrated how social categories can be applied to infer topics one could talk about, during a first encounter. A benefit of social categories and stereotypical information is the minimal amount of information that is needed about the other in order to come up with a first impression.

The core of our Person Memory has been implemented and demonstrated in a small scenario. Yet, a lot of parameters allow for a fine tuning of the system. Determining these parameters is a tedious task. A thorough interaction study is needed in order to identify an optimal set of suitable parameters. Furthermore, we only demonstrated a small subset of the capabilities of our Person Model. Since we focused on first encounters of people, no information about the relationship between agent and interlocutor is available. But behavior towards others is strongly influenced by how good we can relate to each other. E.g., as stated above, the choice of topics during Small Talk is influenced by the relationship of the interlocutors. Therefore, as a next step it has to be investigated how the information of a first encounter, as presented here, can be used in further interactions. The integration of relationship information in equation (2) will e.g. further enhance topic selection.

To increase believability, emotions of the agent should be adapted to the situation, as well. As suggested in section 3.1, behavior towards unknown people should differ from behavior towards people the agent already met. Information of previous interactions, like attitude of the interlocutor towards the agent, will allow altering the agent's emotions and mood. Awareness of what was talked about in previous encounters will help to prevent that the conversation will get annoying over time: It will allow picking up topics of interest of both interlocutors and prevent repetition of already said things.

To conclude, the Person Memory presented in this work provides a solid foundation for further research on human-agent long-term relationships. It enhances the agent's awareness of persons he interacted with and allows the agent to react to individuals in a more human-like fashion.

REFERENCES

- Adomavicius, G. and Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering*, pages 734–749.
- Bickmore, T. and Cassell, J. (2001). Relational Agents: A Model and Implementation of Building User Trust. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '01, pages 396–403, New York, NY, USA. ACM.
- Brom, C., Pešková, K., and Lukavský, J. (2007). What Does Your Actor Remember? Towards Characters with a Full Episodic Memory. In *Proceedings of the 4th ICVS, ICVS'07*, pages 89–101, Berlin, Heidelberg. Springer-Verlag.
- Campbell, R., Grimshaw, M., and Green, G. (2009). Relational agents: A critical review. *The Open Virtual Reality Journal*, 11(7).
- Conway, M. A. (1987). Verifying autobiographical facts. *Cognition*, 26(1):39–58.
- Dahlgren, K. (1985). The Cognitive Structure of Social Categories. *Cognitive Science*, 9(3):379–398.
- Endrass, B., Rehm, M., and André, E. (2011). Planning Small Talk behavior with cultural influences for multiagent systems. *Computer Speech & Language*, 25(2):158–174.
- Hamilton, D. (1979). A cognitive-attribitional analysis of stereotyping. *Advances in experimental social psychology*, 12:53–84.
- Huber, M. (1999). JAM: A BDI-theoretic Mobile Agent Architecture. In *Proceedings of the third annual conference on Autonomous Agents*, pages 236–243. ACM.
- Kasap, Z., Ben Moussa, M., Chaudhuri, P., and Magnenat-Thalmann, N. (2009). Making Them Remember – Emotional Virtual Characters with Memory. *Computer Graphics and Applications, IEEE*, 29(2):20–29.
- Kobsa, A. (2007). Generic User Modeling Systems. In *The adaptive web*, pages 136–154. Springer-Verlag.
- Lessmann, N., Kopp, S., and Wachsmuth, I. (2006). Situated interaction with a virtual human - perception, action, and cognition. In Rickheit, G. and Wachsmuth, I., editors, *Situated Communication*, pages 287–323. Mouton de Gruyter, Berlin.
- Rehm, M., André, E., Bee, N., Endrass, B., Wissner, M., Nakano, Y., Nishida, T., and Huang, H. (2007). The CUBE-G Approach - Coaching Culture-Specific Nonverbal Behavior by Virtual Agents. In *Proceedings of the 38th Conference of ISAGA*.
- Rich, E. (1979). User Modeling via Stereotypes. *Cognitive science*, 3(4):329–354.
- Schneider, K. (1988). *Small Talk: Analysing Phatic Discourse*. Hitzeroth, Marburg.
- Schulman, D., Sharma, M., and Bickmore, T. (2008). The Identification of Users by Relational Agents. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 1*, pages 105–111.
- Senay, I. and Keysar, B. (2009). Keeping Track of Speaker's Perspective: The Role of Social Identity. *Discourse Processes*, 46(5):401–425.
- Stangor, C. and Lange, J. (1994). Mental representations of social groups: Advances in understanding stereotypes and stereotyping. *Advances in experimental social psychology*, 26:357–357.