CHAPTER

27 Affect Elicitation for Affective Computing

Jacqueline M. Kory *and* Sidney K. D'Mello

Abstract

The ability to reliably and ethically elicit affective states in the laboratory is critical in studying and developing systems that can detect, interpret, and adapt to human affect. Many methods for eliciting emotions have been developed. In general, they involve presenting a stimulus to evoke a response from one or more emotion response systems. The nature of the stimulus varies widely. Passive methods include the presentation of emotional images, film clips, and music. Active methods can involve social or dyadic interactions with other people or behavioral manipulation in which an individual is instructed to adopt facial expressions, postures, or other emotionally relevant behaviors. This chapter discusses exemplar methods of each type, discusses advantages and disadvantages of each method, and briefly summarizes some additional methods.

Key Words: affect elicitation, emotional images, emotional film clips, emotional music, backward masking, behavior manipulation, social interaction, dyadic interaction

Introduction

One important goal of affective computing (AC) is to develop computational systems that can recognize and respond to the affective states of the user, thus enabling machines to interact with humans in what is hoped to be a more effective, naturalistic fashion (Picard, 1997). The success of affect-sensitive interfaces depends on their accuracy at detecting or recognizing users' affective states. Affect detectors typically rely on supervised classification techniques to develop models that associate particular behaviors (such as physiological measures and facial features) with emotion labels or dimensions. These detectors need large data sets for model development and validation. As such, it is critical to obtain reliable datasets of emotional expressions collected from a wide range of individuals. These data fall in three primary categories: (1) acted or posed expressions obtained by asking individuals, often actors, to portray emotions; (2) induced expressions, in which emotional responses are elicited via some stimulus; and (3) naturalistic displays of emotion. Acted data are easy to collect, but an important concern is of the data's ecological validity. For example, when told to act frustrated, individuals do not smile, but during natural experiences of frustration, the majority of individuals *do* smile (Hoque & Picard, 2011). Naturalistic data, on the other hand, are certainly ecologically valid but notoriously difficult to collect. The middle road, in which specific emotions are elicited, may be a good compromise for AC researchers and this has been used to great effect already (see, e.g., Bailenson et al., 2008; Khalali & Moradi, 2009; Koelstra et al., 2011; Leon, Clarke, Callaghan, & Sepulveda, 2007; Monkaresi, Hussain, & Calvo, 2012; Soleymani, Pantic, & Pun, 2012). This chapter focuses on method to elicit particular affective states with an emphasis on methods that have been previously used or have considerable potential for use in AC research.

Emotions arise in response to many different stimuli, so a number of diverse strategies for eliciting emotions have been developed. These vary on a range of dimensions, including sensory modality, reliability, intensity of response, temporal length of response, and ecological validity. The most common emotion elicitation methods can be divided into two main categories: (1) passive or perception-based and (2) active or expression-based. In passive methods, individuals observe stimuli—such as film clips, images, or music—that are designed to evoke particular feelings and moods. In active methods, individuals are instructed to perform particular behaviors that might naturally evoke different emotions, such as posing facial muscles, adopting body postures, or interacting with other people. This chapter discusses exemplar methods from each of these categories in ~~Sections 2~~ and ~~3~~. ~~Section 4~~ presents a summary of some less common but still potentially useful methods. Many of the methods and some of the examples discussed in this chapter have been culled from the *Handbook of Emotion Elicitation and Assessment* (2007), and a wealth of additional information can be found in this handbook and in the references for this chapter.

Passive Methods

Emotional Images

In this method, images that evoke particular affective states are presented to an individual, one image at a time. For example, an image of a mutilated person might be presented to elicit disgust, while an image of a family might be used to elicit happiness. The images should have the capability to consistently evoke a particular level of a single affective state in the viewer (e.g., mild pleasure, strong disgust). The presentation method is standardized such that all individuals have the same viewing experience. For example, the images could each be presented for 10 seconds on a computer screen that is a fixed distance from where the individual is sitting, with a constant screen resolution, screen brightness, and image size (Monkaresi, Hussain, & Calvo, 2012). The images themselves can be selected on the basis of which emotions should be elicited from a database of standardized images. One of the most prominent databases of this kind is the International Affective Picture System (IAPS), which includes over a thousand images depicting people, objects, and events that have been standardized on the basis of pleasure and arousal (Bradley & Lang, 2007; Lang, Bradley, & Cuthbert, 2005). Another more recently developed database is the Geneva Affective Picture Database (GAPED) (Dan-Glauser & Scherer, 2011).

Several researchers have used the IAPS to elicit emotions for use in training affect detection systems. For example, Khalili and Moradi (2009) focused on eliciting calm, positively excited, and negatively excited emotions. They selected 70 to 150 pictures for each emotion, which were shown to individuals in sets of five pictures, with each picture shown for 2.5 seconds. A self-reporting phase followed each set. They used these data to train unimodal and multimodal emotion detectors using EEG and peripheral physiological signals. Leon et al. (2007) and Monkaresi et al. (2012) followed similar protocols but collected self-reports after every IAPS image instead of after sets of images.

Using images to elicit emotions is advantageous for several reasons. First, images are noninvasive. They share features with actual physical objects and thus are generally good cues for eliciting the same emotions as the objects they represent (Bradley & Lang, 2007). In addition, because images are static cues, some of the difficulties in working with longer, dynamic stimuli are avoided (such as movement, narrative development, and complicated context). Nearly all aspects of the stimuli can be controlled or manipulated, including image size, duration of presentation, brightness, luminosity, spatial frequency, and color. Images are easily accessible via the databases listed above. It is relatively simple to set up a computer to display them, or, alternatively, one need not even use a computer to show the images.

However, the strength of emotions elicited via static images tends to be lower than when they are elicited via film or in anticipation of real events (discussed below). Emotional reactions to images tend to be short and transient, being present during and shortly after viewing and decaying rapidly thereafter. Physiological reactions to viewing unpleasant photos tend to be low, likely because pictures are only symbolic of people, objects, and events—they are not those real stimuli. As such, images may not be the best choice when a strong emotional reaction is desired. This lessened emotional response, does, however, alleviate concerns about causing participants undue stress when they are viewing particularly unpleasant photos. Bradley and Lang (2007) also suggest that there is no clear one-to-one relation between image content and discrete emotions because the background knowledge and life experience each person has when reacting to or appraising any given image ultimately determines his or her emotional reaction. However, researchers such

as Asensio et al. (2010) have used this fact to their advantage, studying how individuals with cocaine addictions react differently to pleasant, unpleasant, and neutral IAPS images as compared with individuals without an addiction. Another concern is that the individual's mood and the experimental conditions, beyond reaction and habituation times, are generally not accounted for. As such, images might be more useful for inducing dimensional affect (e.g., positive, neutral, negative) rather than discrete affect (e.g., anger, surprise). That said, other research has indicated that IAPS images generally do, in fact, elicit one discrete emotion more than others (Mikels, Fredrickson, Larkin, Lindberg, Maglio, & Reuter-Lornez, 2005). Finally, utilizing images to elicit emotions may be useful for acquiring data from reactive modalities such as facial expressions, physiological signals (e.g., heart rate and skin conductance), and neural signals (e.g., functional near-infrared spectroscopy (fNIRS) or EEG data). However, because of the stationary visual nature of image presentation, images may not be as suitable for obtaining data from productive modalities, such as acoustic-prosodic cues, text or discourse, or gestures.

Emotional Film Clips

Short films can be presented to elicit target emotions (Rottenberg, Ray, & Gross, 2007). For comparison purposes, a neutral baseline film is generally shown prior to the presentation of each emotional film clip. For example, Soleymani et al. (2012) showed participants a series of 20 film clips to elicit emotions ranging from calm, aroused, pleasant, unpleasant, and neutral emotions. In each trial, a 15-second neutral clip was shown, then a 1- to 2-minute emotional clip, followed by a self-assessment phase. Video, sound, eye gaze, EEG, and various physiological signals were collected. Soleymani et al. (2012) used the eye gaze and EEG data to train and validate unimodal and multimodal affect classifiers. Instead of using short clips, Bailenson et al. (2008) recorded physiological data and video while individuals viewed a single 9-minute clip that contained amusing, neutral, and sad sections. There was no self-report of affect phase. Instead, trained coders rated the videos of individuals' facial expressions. The data was used to build various affect classifiers.

In presenting films as in presenting images, the physical situation should be standardized. For example, Rottenberg et al. (2007) suggest showing the films on a 20-inch monitor about 5 feet from the participant. Then, for each target emotion, one or two short clips that are 1 to 3 minutes in length and that are as homogenous as possible should be shown. All the film clips used should be matched on relevant dimensions so that when multiple emotion film conditions are compared, researchers can infer that any effects seen are due to the emotions, not differences in the films. Relevant matching characteristics include the length of the clip, its intensity, complexity, core themes, presence and number of humans, color, brightness, picture motion, and theoretically important dimensions such as activation level (see Rottenberg et al., 2007; also see Detenber, Simons, & Bennet, 1998). However, it is rarely possible to match films on every dimension.

Rottenberg et al. (2007) presented a list of film clip recommendations for eliciting seven different emotions (amusement, anger, disgust, fear, neutral, sadness, and surprise), with self-reported emotions for each clip and instructions for obtaining each clip. A larger database of emotion-eliciting films is discussed by Schaefer, Nils, Sanchez, and Philippot (2010), with many film clips available online (http://www.ipsp.ucl.ac.be/recherche/FilmStim/). For example, an abstract visual display from a screensaver can work well as a baseline/neutral film (Gross & Levenson, 1995). Another good baseline is a pleasant/low contentment clip depicting nature scenery, animals, and uplifting music, such as a segment from the nature film *Alaska's Wild Denali* (Hardesty, 1997). Clips depicting death scenes from the movies *The Champ* (Lovell & Zeffirelli, 1979) and *The Lion King* (Hahn, Allers, & Minkoff, 1994) can be used to elicit sadness, while clips from comedies such as *When Harry Met Sally* (Reiner, Scheinman, Stolt, & Nicolaides, 1989) and *Bill Cosby, Himself* (Cosby, 1996) can elicit amusement.

Films are advantageous because they capture attention well. They can elicit higher intensity emotions and more complex emotional states than images (Rottenberg et al., 2007). As such, they may be a good choice for studying the magnitude of emotional responses. Film clips may also be useful for studying emotion latency, rise time, duration, and offset because—unlike still images—they are dynamic stimuli, occurring as a sequence of images and sound through time.

It should be noted that the ecological validity of emotion displayed in film clips is in question. On one hand, the affectively charged situations depicted in films appear to be real. However, as Rottenberg et al. (2007) discuss, emotional responses to films require the willing suspension of disbelief. Films,

like still images, are a step removed from reality. Participants may react emotionally in spite of *or* because of this as a result. An additional problem is that individuals may have previously seen the films from which the clips are taken. A film clip may not have as much impact if an individual has seen the whole film, or it may have a much stronger impact, since the individual knows the whole context of the clip. Finally, as could be seen in the examples above, film clips are useful for collecting facial expressions, physiological signals, and neural signals. However, like images, they are less ideal for collecting gestures or verbal behavior.

Music and Ideation

Music has generally been used to induce moods in psychological research. A mood can be defined as an enduring affective state that may not directed toward a specific object or event—that is, it is an underlying feeling versus an emotional episode, such as fear or surprise, that is directed at an object or event, such as a snake or a surprise party (Russell, 2003; Watson, 2000). However, Russell (2003) notes that the duration and stability of a mood is not well defined. Music, like moods, is also less object- or event-focused, unlike films or images. When music is combined with a secondary mood-induction technique, robust results can be obtained, such as a higher rate of successfully induced moods, stronger induced moods, and moods that are stable over time and across tasks (Eich, Ng, Macaulay, Percy, & Grebneva, 2007; Vastfjall, 2002).

In one study, Kim, Bang, and Kim (2004) used a combination of background music, a narrated story, and lighting color changes to elicit sadness, anger, stress, and surprise. For example, sadness might be elicited by using slower music, a sad story, and blue light, while the use of music with a faster tempo and red light might elicit anger. Each participant was presented with the four scenarios, with each followed by a self-report. Kim et al. (2004) collected physiological data, which they used to develop affect classifiers. Koelstra et al. (2011) showed music videos to participants to collect similar data.

Eich et al. (2007) describe a technique that combines music with contemplation/ideation and an idiographic approach (abbreviated MCI). Individuals listen to selections of cheerful or melancholy music while contemplating happy or depressing thoughts. Only after an individual is known to be sufficiently in the target mood state (via periodic self-reports) do further experimental procedures occur. Unlike some other music and ideation procedures, the MCI technique ensures that all individuals reach the same predetermined level of the target mood, although the length of the mood-induction period will likely vary for each individual (see Vastfjall, 2002, for a comprehensive review of musical mood induction procedures; see Westermann, Spies, Stahl, & Hesse, 1996, for a review of other mood induction procedures). Eich et al. (2007) note that about 80% of the participants in their studies were able to develop the desired mood (very pleasant or very unpleasant) and that the mood elicited tended to be reasonably strong, stable, sincere, and reproducible.

The musical pieces used by Eich and colleagues (Eich et al., 2007; Eich & Metcalfe, 1989; Ryan & Eich, 2000) in the MCI were originally selected based on educated guesses. They are all classical compositions ranging in duration from 2 to 6 minutes. A full list is presented in Eich et al. (2007). The happy music includes several livelier pieces such as Vivaldi's *Four Seasons:* Spring I Allegro and Spring III Allegro, Mozart's *Eine Kleine Nachtmusik:* Allegro, and Tchaikovsky's *The Nutcracker:* Waltz of the Flowers. The sad music samples a wider set of composers, including Albinoni's Adagio in G Minor, Chopin's Prelude #4 in E Minor, and an excerpt from Stravinsky's *Firebird*: Lullaby. Eich et al. (2007) also point out several ways to improve music selection. For example, researchers could draw on the knowledge of musicology experts about which elements of music are key to inducing particular emotions. More care could be taken to match styles and genres of the music used. Self-chosen music is also an option, since some people may associate particular songs with the moods that are to be induced, and these songs may work better for the induction (e.g., see Carter, Wilson, Laweson, & Bulik, 1995; Vuoskoski & Eerola, 2012).

One disadvantage of using the MCI is that it takes time—on average, Eich et al. (2007) found that it takes individuals 15 to 20 minutes to reach an adequate level of pleasure or displeasure. However, it is conceivable that an individual's affective state could be changed in a shorter time. Songs are often 3 to 4 minutes in length, and musicians attempt to draw the listener through one or more states during the song. A second concern is that individuals are often explicitly told that the researchers are trying to get them into a positive or negative mood. This means that the individuals will possibly merely act in a manner consistent with their beliefs about moods rather than feeling the intended mood. To allay these concerns, Eich et al. (2007) recommend

collecting self-reports from individuals about how genuine they think their moods are. Previously, individuals have generally reported that they do think their moods are genuine (Eich & Metcalfe, 1989; Ryan & Eich, 2000).

Music seems to be most useful for collecting reactive emotional expressions, such as physiological and neural signals. Because of its passive nature—the participant must sit still and listen to the music—it may not lend itself as well to data collection from other modalities.

Backward Masking and the Dissociation Paradigm

The previous three methods attempt to elicit emotion through the presentation of particular kinds of stimuli (such as films or music). This next method, however, studies a different kind of emotional response altogether: unconscious emotions. It has been argued that in emotional processing, stimuli are evaluated or appraised by both conscious and unconscious mechanisms (Berridge & Winkielman, 2003; Kihlstrom, 1999; LeDoux, 1996; Öhman, 1986; Wiens & Öhman, 2007). The presence of unconscious emotion can be inferred from changes in emotional experience, behavior, and psychophysiology if either (1) the stimulus was manipulated so that conscious awareness of the stimulus was prevented, such that any emotional response would be due to unconscious rather than conscious mechanisms, or (2) individuals display signs of emotional processing without reporting changes in their (conscious) emotional experience.

Backward masking in the *dissociation paradigm* is the most common method for manipulating the stimulus to elicit unconscious emotions (Wiens & Öhman, 2007). Conscious awareness of a visual stimulus is blocked in two steps: (1) the stimulus (also called the target) is shown very briefly, as for 15 to 60 milliseconds, then (2) a second visual stimulus (also called the mask) is shown for a longer duration, such as 500 milliseconds. Individuals tend to report being consciously aware only of the mask. Other masking techniques include sandwich masking (mask-target-mask) and energy masking (the mask is a light flash). The dissociation paradigm is an experimental design in which awareness of the target must be eliminated in order to test whether this awareness is necessary for responding. If an individual is unaware of the target, then any response to the masked target is the result of unconscious processes rather than the direct involvement of conscious awareness of the target. For this paradigm to succeed, it must be demonstrated that the participant actually *is* unaware of the target. Furthermore, it is assumed that awareness does play some causal role in determining a participant's responses and is not epiphenomenal to emotional processing. Wiens and Öhman (2007) discuss in detail how unawareness can be measured; also see Cheesman and Merikle, 1984; Merikle and Reingold, 1998; and Öhman and Wiens, (1994).

The main practical issue to consider is stimulus presentation. Masking a stimulus requires (1) a very fast change from the first visual display to the next and (2) a stable picture presentation time across trials to avoid confounds. Two commonly used displays are cathode-ray tube (CRT) monitors and thin-film transistor (TFT) or liquid-crystal display (LCD) screens. Using these, however, limits the minimum stimulus display time owing to screen refresh rates. Wiens and Öhman (2007) recommend placing a mechanical shutter in front of a data projector that does not rewrite the screen after each refresh cycle (i.e., those based on TFT or LCD technology), thus allowing greater accuracy of presentation durations as well as millisecond-level control.

The method described here has been used successfully several times in the affective sciences but rarely by AC researchers. In a classic study, Öhman and Soares (1994) studied unconscious fear processing via a forced-choice classification task. On each trial, participants were shown a masked picture of a spider, snake, flower, or mushroom and had to indicate which of these four items the picture showed. Skin conductance and emotional ratings were recorded to index emotional processing. Öhman and Soares (1994) found that participants who were spider-phobic showed greater skin conductance responses to spiders but not snakes, while participants who were snake-phobic responded to snakes but not spiders. Neutral participants did not respond to either. This suggested that fear could be present without a person being consciously aware of the feared stimulus.

In another study, Winkielman, Berridge, and Wilbarger (2005) showed masked emotional faces to participants who were either thirsty or not thirsty. The participants then had to rate their mood and were allowed to consume a fruity drink. Winkielman et al. (2005) found that the consumption behavior of the thirsty participants changed in relation to the emotion of the face they had seen, although no mood differences were reported across these participants. Again, this seems to indicate that

an emotional response can occur even without an effect on conscious emotional experience.

Although this method has not seen much if any use in AC research, it is still a potentially important technique if the goal is to train models to detect subtle unconscious emotional responses. A related method that has seen use places the visual stimulus in the background rather than masking it entirely or, as described in ~~Section 2.1~~, placing it in the foreground. Hussain, Calvo, and Chen (2013) placed images behind math tasks in order to study the interactions between emotion and cognitive load.

Active Methods

Behavioral Manipulation

Emotional episodes can generally include three components: (1) an eliciting event or stimulus, (2) an emotional feeling (for conscious emotions), and (3) emotional behaviors, such as facial expressions, physiological responses, and physical actions. The passive methods discussed so far have focused on (1)—that is, these methods vary the stimuli an individual encounters in order to evoke different emotions in that individual. The active methods we discuss next focus instead on (3), wherein individuals are instructed to adopt particular behaviors or expressions in order to change the emotions they experience. The first of these involves behavioral manipulation.

In this method, individuals are instructed to adopt particular muscular configurations or behavioral patterns that have been associated with emotional experiences, such as contracting or relaxing facial muscles or exaggerating natural emotional expressions (Ekman & Davidson, 1993; Laird, 1974; Levenson & Ekman, 2002). This, as Ekman (2007) argues, will generate the same kinds of automatic nervous system (ANS) activity patterns as are generated by emotional experiences. Ekman and colleagues developed a well-known method for eliciting emotion through the manipulation of facial expressions called the directed facial action task (Ekman, 2007; Ekman & Davidson, 1993; Levenson & Ekman, 2002). In this task, individuals are given general instructions about how to voluntarily move their facial muscles in particular ways in order to produce emotional responses. They are *not* told to pose a particular emotion. One example of this method in AC research can be found in Vural et al. (2007), who used a database of directed facial actions to develop a facial expression classifier.

In other muscle-by-muscle manipulations, researchers have tried to disguise the instructions given to participants in order to see whether muscle movement alone is sufficient to elicit emotion. For example, Zajonc, Murphy, and Inglehart (1989) had English-speaking participants pronounce various sounds as a way of inducing facial expressions. Strack, Martin, and Stepper (1988) recruited participants for a study purportedly about teaching writing to handicapped people. The participants were instructed to hold a pencil in their mouths with either their lips tightly clamped (as in a smile) or with their lips drawn back (as in an expression of disgust). After performing several writing exercises with the pencils in their mouths, participants rated the funniness of cartoons. Strack et al. (1988) found that participants who held the pencil in a smile-like expression rated the cartoons as funnier than those who had a disgust-like expression.

Another strategy is to modify existing emotional expressions, rather than attempting to create them (see Laird & Strout, 2007, for detailed discussion of this methodology). For example, Lanzetta, Cartwright-Smith, and Kleck (1976) asked participants to either inhibit or intensify their expressive reactions to a series of uncomfortable electric shocks. Lanzetta et al. (1976) recorded skin conductance as well as self-reports. When they were exaggerating their reactions, participants had higher skin conductance and reported that the shocks were more painful. Vice versa, when participants inhibited their reactions, they reported the shocks as being less painful and had lower skin conductance. In an earlier study by Bandler, Madaras, and Bem (1968), participants were instructed to place their hands on plate to receive an electric shock. One group was instructed to leave their hands on the plate; the other group was instructed to jerk their hands away quickly. Although the first group endured a longer duration of shock, the second described the shocks as more painful because of the jerk-away action. This highlights how performing emotional actions can affect the experienced emotion. Indeed, Duclos and Laird (2001) have suggested that one effective way of eliciting anger through behavior is to ask individuals to act angry.

Posture can be manipulated in much the same way as facial expressions, using muscle-by-muscle or exaggeration/minimization instructions (Duclos, Laird, Schneider, Sexter, Stern, & Van Lighten, 1989; Flack, Laird, & Cavallaro, 1999). Other behavioral manipulations that have been successful, though with generally weaker results than those mentioned above, include manipulations of breathing (Philippot, Chapelle, & Blairy, 2002),

eye gaze patterns (Schnall et al 2000; Williams & Kleinke, 1993), and tone of voice (Hatfield et al., 1995; Siegman & Boyle, 1993) (see Laird & Strout, 2007, for some further discussion). Combinations of methods, such as expressions and postures (e.g., Flack et al., 1999), can produce more intense effects.

If one knows the physical behaviors associated with the target emotions, behavioral manipulation techniques can be fairly precise in eliciting the target, since behaviors are often distinct. The effects are moderately strong, though the strength of elicited emotional response often still falls short of natural or other induced methods (for example, from films). The effects vary across individuals as well. Some people report being relatively unaffected by posing facial expressions of happiness or anger, while others are dramatically affected. Laird and Strout (2007) point out that researchers using other methods, such as the passive methods described above, have not explored these individual differences as much and thus very likely encounter similar variations. Behavioral manipulations also have the advantage of not containing other cognitive material (as a film might), not containing verbal interaction (unless a verbal behavioral manipulation is being used), not taking too much time, and not involving very elaborate methodologies (Laird & Strout, 2007).

However, there are concerns about ecological validity—is eliciting emotions just from physical behaviors tantamount to inducing the emotions out of their context? Can one really say that emotions elicited through behavioral manipulations are more "pure" (as Laird & Strout, 2007 do) if emotions outside the laboratory occur in the complicated and messy context of the real world? Similar criticisms can be applied to most lab-based elicitation methods. In addition, most behavioral manipulations require the researcher to already have an understanding of which behaviors to perform or how to move one's muscles in ways to induce the target emotion. Although facial expressions linked to Ekman's six "basic" emotions have been identified (Ekman, 2007), relatively less research has looked at distinct facial expressions or behaviors related to, for example, more complex affective states like frustration, confusion, and engagement. This limits the range of possible target emotions that can be elicited with these methods. Furthermore, the poses or actions required can be difficult or complicated for individuals to perform and not all individuals may be able to perform them.

Behavioral manipulations may be most useful for collecting reactive expressions, such as physiological and neural signals, since the productive aspects of the emotional expression are tightly controlled.

Social Psychological Methods and Social Interaction

In this method, which comes from social psychology, researchers try to create realistic social scenarios that elicit emotions in a more naturalistic context (Harmon-Jones, Amodio, & Zinner; 2007). Manipulations tend to be high-impact. Deception is generally needed to keep participants unaware of variable manipulations. For example, to elicit high anxiety, participants in a study may be told that they will be given a series of electric shocks (e.g., Schachter, 1959). To induce anger, a participant may be given insulting evaluations or feedback from a fake "second participant" in the study (e.g., Harmon-Jones & Sigelman, 2001).

Not many examples using social psychological methods are found in the AC literature. There are however some notable exceptions. The SEMAINE database recorded interactions between users and four emotionally stereotyped characters in order to study social signals in conversations between humans and artificial intelligent agents as well as to collect data for training such agents (McKeown, Valstar, Cowie, Pantic, & Schroder, 2012). The characters were designed to draw users into their emotional states—happy, angry, sensible, and depressive. They were played either by a human operator who pretended to be an artificial agent or an autonomous but more limited artificial agent. The database has been used to develop multimodal affect classifiers (e.g., Wöllmer, Kaiser, Eyben, & Schuller, 2012). Kim (2007) had individuals play a structured quiz game with a virtual character in which a hidden person guided the course of the quiz in order to elicit certain emotions. Although this comes closer to the kind of social interaction that Harmon-Jones et al. (2007) discuss, it again is not a high-impact manipulation.

Harmon-Jones et al. (2007) laid out five critical steps for setting up social psychology studies with deception. (1) Construct a cover story as rationale for the experiment to distract participants from the experiment's true purpose. (2) Determine what the experimenter's behavior will be, such that it will be consistent over the data collection period and across participants. (3) Design a manipulation of the independent variable that minimizes participant awareness of the manipulation while ensuring that they respond just as they would if they encountered the stimulus outside the lab.

Make the experimenters blind to the condition when possible. (4) For the dependent variable, measure emotional responses with a behavioral, physiological, or self-report method. (5) Conduct a postexperimental interview to check for clarity of instructions, suspicion about the experiment's true goals, and to debrief the participant about purpose of experiment.

The cover story will vary with the emotion elicited. For example, to elicit joy and sadness, researchers have used social-comparison manipulations. Participants perform a task that is described as easy or difficult, such as solving a set of analogies, in order to induce an initial worry or stress state (e.g., Forgas, Bower, & Moylan, 1990). Then participants are given feedback on how they performed on the task relative to others. Positive feedback, such as being told that they are above average, leads to increased joy, while negative feedback leads to increased sadness.

Sympathy and guilt are two other complex emotions that can be elicited (Harmon-Jones et al., 2007). Sympathy scenarios generally involve perspective taking. For example, the participant may be asked to listen to a radio broadcast about a child diagnosed with cancer and then asked to imagine how the child and his or her family may feel (Harmon-Jones, Peterson, & Vaughn, 2003). Guilt requires more deception to elicit and has been successfully elicited using a false physiological feedback procedure (e.g., Monteith, Ashburn-Nardo, Voils, & Czopp, 2002). Participants who were not prejudiced were told that physiological arousal increases when a person is viewing negative images. Participants were then shown a series of images of people of different races, with each image followed by the supposed physiological measurements taken during its viewing. In reality, the measurements were faked such that participants were manipulated into thinking that they had prejudiced reactions to seeing people of other races, thus eliciting guilt.

The biggest advantage of using social psychological methods is that the emotions elicited through cover stories and social contexts are more realistic and similar to emotions occurring in the real world. Some emotions that are more difficult to elicit with other methods, such as anger and guilt, can be elicited with social psychological methods (Harmon-Jones et al., 2007). In addition, Harmon-Jones et al. (2007) suggest that issues present in other methodologies, such as responses that are merely due to demand characteristics, may be circumvented through social psychological methods. However, more realistic scenarios also tend to allow a wider variability in responding, which can make it harder to elicit the specific discrete emotion one may want to study.

Social interaction thus seems best suited for collecting productive expressions of emotion, including body movements, gestures, acoustic-prosodic cues, and discourse or language. Facial expressions, if cameras are positioned appropriately, can be collected. With mobile equipment, physiological signals can be recorded. Neural signals, however, tend to require participants to be seated and stationary, which would be difficult in a social scenario.

Dyadic Interaction

In this method, a dyad is brought into the lab rather than an individual. The dyad is then instructed to engage in an unrehearsed, minimally structured conversation (Roberts, Tsai, & Cohen, 2007). Dyads can be romantic couples (e.g., Cohan & Bradbury, 1997; Gonzaga, Keltner, Londahl, & Smith, 2001; Tsai & Levenson 1997), siblings (e.g., Shortt & Gottman, 1997), or children and caregivers (e.g., Messinger et al., this volume; Weis & Lovejoy, 2002; Reppeti & Wood, 1997). Sometimes, an unfamiliar dyad is brought in, such as ethnically similar or dissimilar pairs, in order to study other aspects of interpersonal relations (e.g., Littleford, Wright, & Sayoc-Parial, 2005). Like social psychological methods, dyadic interactions have not been used nearly as much as other emotion elicitation methods in AC. One example is a corpus of human-human interactions collected by Zara, Maffiolo, Martin, and Devillers (2007) to aid in the design of affective interactive systems. Pairs of strangers played an emotional adaptation of the game Taboo, in which one player has to guess a word that the other player describes through speech and gesture but without using any of five forbidden words.

Dyadic interactions draw on the social context of emotion to elicit a wide range of spontaneous and realistic emotions and follow their natural temporal course while still remaining in a controlled situation. Therefore it has high ecological validity, perhaps higher than any other elicitation methods. Emotions often occur in social contexts and have social functions. Because the dyad regularly interacts outside the lab or is a pair of strangers, there is not the same worry as in other social interaction methods that reactions will be less natural due to the laboratory context. Emotions are dynamic, and this kind of scenario allows one to study their

time course. A range of emotional responses can be elicited by varying (1) the kind of conversation the dyad can have and (2) individual factors that change the intensity of emotion displays. For example, Coan, Gottmann, Babcock, and Jacobson (1997) asked couples to discuss disagreements in their relationship in order to elicit intense negative emotion. Similarly, discussions of enjoyable topics can be used to elicit positive emotion (Roberts et al., 2007). The interaction is usually synchronous, but does not have to be—one recent study elicited emotions through an asynchronous dyadic interaction wherein one individual provided written feedback on another individual's essay (Calvo, personal communication, May 27, 2013).

However, the minimal structure of this method allows room for participant noncompliance and experimenter error. There can be great variability in emotional responding, and the facilitator for the task may establish different rapport with each dyad or otherwise act differently, which could influence emotional responses. Therefore this method may not be as well suited for eliciting precise, discrete emotions as it is for eliciting natural emotions over time. Because tasks are often 2 to 4 hours in length, significant resources are required (Roberts et al., 2007). In addition, the dyadic interaction provides only a snapshot sampling of emotion. Researchers lack the history of the dyad's relationship and have limited knowledge of why each member of the dyad responds to the other in the ways they do. That said, for some relationships (e.g., teacher-student or parent-child), some assumptions about the relationship can be made (e.g., power dynamics).

Dyadic interactions seem to be useful for the collection of much the same data as social interactions—gestures, body movements, verbal behavior, facial expressions, and physiological data.

Other Methods

This chapter has provided an introduction to the wide range of emotion elicitation methods available. We selected several of the most promising and popular methods to discuss, but there are many more. For example, individuals can perform a writing task, such as the autobiographical memory task (Dunn & Schweitzer, 2005; Myers & Tingley, 2001), in which they are prompted to recall and write about an emotional event. Stevenson and James (2008) have proposed an International Affective Digitized Sounds (IADS) database of affective auditory stimuli. Other researchers have instructed individuals to read and internalize positively or negatively valenced self-referential statements (Seibert & Ellis, 1991; Velten, 1968), used stimulus-reinforcement association learning to elicit emotions (Rolls, 2007) or have hypnotized individuals to make them think emotionally provocative thoughts (Friswell & McConkey, 1989). Emotional functioning can also be studied in neurologically impaired patients via lesion studies, brain imaging studies, and neural activation studies (for a review, see Levenson, 2007).

Conclusions

Many methods exist for eliciting emotions. Passive methods tend to be less ecologically valid and elicit less naturalistic emotions. They can, however, be easier to implement. In addition, most aspects of the stimuli—such as the color, size, and duration of presentation for an image—can be controlled with passive methods. Emotions elicited through passive methods vary in intensity and complexity, from less intense via images to more intense films or music. Active methods have a similar range of intensity and complexity: behavior manipulations are very precise but may not produce very strong responses, while social and dyadic interactions can elicit strong naturalistic emotions but are more imprecise. Active methods tend to have higher ecological validity and can elicit complex emotions, such as anger and guilt, that may be hard to achieve via other methods but may also have a wider variability in responses.

The stimuli used in contemporary emotion elicitation methods can be temporally short, such as images, or stretched through time, such as films or social interactions. As such, care must be taken when selecting methods for collecting participants' emotional responses to serve as manipulation checks that the intended emotion was in fact induced. Rottenberg et al. (2007) note that self-report questionnaires, which are often administered immediately following the stimuli, have two main problems. First, delays between the activation of an emotion and the assessment of the emotion can introduce measurement error. Second, collecting only retrospective emotion ratings has very low temporal resolution. To get finer-grained measurements, Rottenberg et al. (2007) suggest using a continuous measure of emotion experience, such as a rating dial method (validated by Levenson & Gottman 1983; also see Gottman & Levenson, 1985), recording video of expressive behavior during stimuli presentation, or recording central and peripheral physiological responses during stimuli presentation. The selection of appropriate ways of measuring emotion is just as crucial as the choice

of method of emotion elicitation itself. Another concern is to ensure that an emotion elicitation method does not inadvertently induce nonintended emotions in addition to the intended emotion. For example, a method to induce *fear,* such as a clip of a traffic accident, might also induce *sadness* if a child is seriously injured. In fact, preventing the induction of untended emotions or mixed emotions is one of the most serious challenges facing most emotion elicitation methods.

Finally, as mentioned in the Introduction, emotion elicitation is only one of several methods for collecting emotional data for AC research. Where ecological validity is concerned, there is a clear hierarchy: recording natural emotional expressions is best, acted data less so, and elicited emotion varying in the middle. Training data have an important influence on affect classification. In a recent meta-analysis of 30 multimodal affect classifiers, the performance improvements of the multimodal classifiers over the best unimodal classifiers was three times higher (12.1% improvement) when trained on acted data as opposed to natural or elicited data (4.39% improvement) (D'Mello & Kory, 2012). The classifiers trained on natural or elicited data may better reflect current affect detection capabilities, since the ultimate goal of AC is to build interfaces that sense and respond to naturalistic expressions of emotion in the real world. Indeed, it seems prudent to use data that reflect, as closely as possible, the situations and emotional expressions that these interfaces will encounter—that is, naturalistic data. Owing to the difficulties of collecting completely natural affective data, in this chapter we surveyed what could be considered a good middle road for AC researchers: methods for emotion elicitation.

Acknowledgments

This research was supported by the National Science Foundation (NSF) (ITR 0325428, HCC 0834847, DRL 1235958, and NSF Graduate Research Fellowship under 1122374). Any opinions, findings and conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSF.

References

Asensio, S., Romero, M. J., Palau, C., Sanchez, A., Senabre, I., Morales, J. L.,…Romero, F. J. (2010). Altered neural response of the appetitive emotional system in cocaine addiction: an fMRI Study. *Addiction Biology*, *15*(4), 504–516.

Bailenson, J. N., Pontikakis, E. D., Mauss, I. B., Gross, J. J., Jabon, M. E., Hutcherson, C. A.,…John, O. (2008). Real-time classification of evoked emotions using facial feature tracking and physiological responses. *International Journal of Human-Computer Studies*, *66*(5), 303–317.

Bandler, R. J., Madaras, G. R., & Bem, D. J. (1968). Self observation as a source of pain perception. *Journal of Personality and Social Psychology*, *9*, 205–209.

Berridge, K. C., & Winkielman, P. (2003). What is an unconscious emotion: The case for unconscious "liking." *Cognition & Emotion*, *17*, 181–211.

Bradley, M. M., & Lang, P. J. (2007). The International Affective Picture System (IAPS) in the study of emotion and attention. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of Emotion Elicitation and Assessment* (pp. 29–46). New York, NY: Oxford University Press.

Carter, F. A., Wilson, J. S., Lawson, R. H., & Bulik, C. M. (1995). Mood induction procedure: Importance of individualizing music. *Behaviour Change*, *12*, 159–161.

Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology*, *40*, 343–367.

Coan, J., & Allen, J. J. B. (Eds.) (2007). *Handbook of emotion elicitation and assessment*. New York, NY: Oxford University Press.

Coan, J., Gottman, J. M., Babcock, J., & Jacobson, N. (1997). Battering and the male rejection of influence from women. *Aggressive Behavior*, *23*(5), 375–388.

Cohan, C. L., & Bradbury, T. N. (1997). Negative life events, marital interaction, and the longitudinal course of newlywed marriage. *Journal of Personality and Social Psychology*, *73*(1), 114–128.

Cosby, W. (Producer/Director). (1996). *Bill Cosby, himself* [Motion picture]. United States: Twentieth Century Fox.

D'Mello, S., & Graesser, A. C. (2010). Multimodal semi-automated affect detection from conversational cues, gross body language, and facial features. *User Modeling and User-Adapted Interaction*. *20*(2), 147–187.

D'Mello, S. K., & Kory, J. (2012). Consistent but modest: A meta-analysis on unimodal and multimodal affect detection accuracies from 30 studies. In L. P. Morency, D. Bohus, & H. Aghajan (Eds.), *Proceedings of the 14th ACM International Conference on Multimodal Interaction*, (pp. 31–38). New York, NY: Association for Computing Machinery.

Dan-Glauser, E. S., & Scherer, K. R. (2011). The Geneva affective picture database (GAPED): A new 730-picture database focusing on valence and normative significance. *Behavior Research Methods*, *43*(2), 468–477.

Detenber, B. H., Simons, R. F., & Bennet, G. G. (1998). Roll 'em! The effects of picture motion on emotion responses. *Journal of Broadcasting and Electronic Media*, *21*, 112–126.

Duclos, S. E., & Laird, J. D. (2001). The deliberate control of emotional experience through control of expressions. *Cognition and Emotion*, *15*, 27–56.

Duclos, S. E., Laird, J. D., Schneider, E., Sexter, M, Stern, L., & Van Lighten, O. (1989). Emotion-specific effects of facial expressions and postures on emotional experience. *Journal of Personality and Social Psychology*, *57*, 100–108.

Dunn, J. R., & Schweitzer, M. E. (2005). Feeling and believing: The influence of emotion on trust. *Journal of Personality and Social Psychology*, *88*(5), 736–748.

Eich, E., & Metcalfe, J. (1989). Mood dependent memory for internal versus external events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*, 443–455.

Eich, E., Ng, J. T. W., Macaulay, D., Percy, A., & Grebneva, I. (2007). Combining music with thought to change mood. In

J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp. 124–136). New York, NY: Oxford University Press.

Ekman, P. (2007). The directed facial action task: Emotion responses without appraisal. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp. 47–53). New York, NY: Oxford University Press.

Ekman, P., & Davidson, R. J. (1993). Voluntary smiling changes region brain activity. *Psychological Science*, *4*, 342–345.

Flack, W. F., Laird, J. D., & Cavallaro, L. A. (1999). Additive effects of facial expressions and postures on emotional feelings. *European Journal of Social Psychology*, *29*, 203–217.

Forgas, J. P., Bower, G. H., & Moylan, S. J. (1990). Praise or blame? Affective influences on attributions for achievement. *Journal of Personality and Social Psychology*, *59*(4), 809–819.

Friswell, R., & McConkey, K. M. (1989). Hypnotically induced mood. *Cognition & Emotion*, *3*(1), 1–26.

Gonzaga, G. C., Keltner, D., Londahl, D. A., & Smith, M. D. (2001). Love and the commitment problem in romantic relations and friendship. *Journal of Personality and Social Psychology*, *81*(2), 247–262.

Gottman, J. M., & Levenson, R. W. (1985). A valid procedure for obtaining self-report of affect in marital interaction. *Journal of Consulting and Clinical Psychology*, *53*, 151–160.

Gross, J. J., & Levenson, R. W. (1995). Emotion elicitation using films. *Cognition & Emotion*, *9*(1), 87–108.

Hahn, D. (Producer), Allers, R., & Minkoff, R. (Directors). (1994). *The lion king* [Motion picture]. Burbank, CA: Walt Disney Pictures.

Hardesty, T. (Producer). *Alaska's Wild Denali: Summer in Denali National Park* [Motion picture]. Anchorage, AK: Alaska Video Postcards.

Harmon-Jones, E., Amodio, D. M., & Zinner, L. R. (2007). Social psychological methods of emotion elicitation. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp. 91–105). New York, NY: Oxford University Press.

Harmon-Jones, E., & Sigelman, J. (2001). State anger and prefrontal brain activity: Evidence that insult-related relative left prefrontal activation is associated with experienced anger and aggression. *Journal of Personality and Social Psychology*, *80*, 797, 803.

Harmon-Jones, E., Peterson, H., & Vaughn, K. (2003). The dissonance-inducing effects of an inconsistency between experienced empathy and knowledge of past failures to help: Support for the action-based model of dissonance. *Basic and Applied Social Psychology*, *25*, 69–78.

Hatfield, E., Hsee, C. K., Costello, J., Weisman, M. S., & Denney, C. (1995). The impact of vocal feedback on emotional experience and expression. *Journal of Social Behavior and Personality*, *10*(2), 293–312.

Hoque, M., & Picard, R. W. (March, 2011). Acted vs. natural frustration and delight: Many people smile in natural frustration. In *Automatic face & gesture recognition and workshops (FG 2011)*, 2011 IEEE International Conference (pp. 354–359). New York: Institute of Electrical and Electronics Engineers.

Hussain, S., Calvo, R. A., Chen, F. (2013). Automatic cognitive load detection from face, physiology, task performance and fusion during affective interference. *Interacting with Computers*, first published online June 6, 2013.

Kihlstrom, J F. (1999). The psychological unconscious. In L. A. Pervin & O. P. John (Eds.), *Handbook of personality: Theory and research* (2nd ed., pp. 424–442). New York: Guilford Press.

Kim, J. (2007). Bimodal emotion recognition using speech and physiological changes. In M. Grimm and K. Kroschel (Eds.), *Robust speech recognition and understanding* (pp. 265–280), Vienna, Austria: I-Tech.

Kim, K. H., Bang, S. W., & Kim, S. R. (2004). Emotion recognition system using short-term monitoring of physiological signals. *Medical and biological engineering and computing*, *42*(3), 419–427.

Khalali, Z., and Moradi, M. (2009). Emotion recognition system using brain and peripheral signals: Using correlation dimension to improve the results of EEG. In *Proceedings of the International Joint Conference on Neural Networks* (pp. 1571–1575). New York: Institute of Electrical and Electronics Engineers.

Koelstra, S., Muhl, C., Soleymani, M., Lee, J. S., Yazdani, A., Ebrahimi, T., . . . Patras, I. (2012). DEAP: A database for emotion analysis using physiological signals. *IEE Transactions on Affective Computing*, *3*(1), 18–31.

Laird, J. D. (1974). Self-attribution of emotion: The effects of expressive behavior on the quality of emotional experience. *Journal of Personality and Social Psychology*, *33*, 475–486.

Laird, J. D., & Strout, S. (2007). Emotional behaviors as emotional stimuli. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp. 54–64). New York, NY: Oxford University Press.

Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (2005). *International Affective Picture System (IAPS): Affective ratings of pictures and instruction manual*. Technical report no. A-6. Gainesville, FL: University of Florida.

Lanzetta, J. T., Cartwright-Smith, J., & Kleck, R. E. (1976). Effects of nonverbal dissimulation on emotional experience and autonomic arousal. *Journal of Personality and Social Psychology*, *33*, 354–370.

LeDoux, J. (1996). *The emotional brain: The mysterious underpinnings of emotional life*. New York: Simon & Schuster.

Leon, E., Clarke, G., Callaghan, V., & Sepulveda, F. (2007). A user-independent real-time emotion recognition system for software agents in domestic environments. *Engineering Applications of Artificial Intelligence*, *20*(3), 337–345.

Levenson, R. W. (2007). Emotion elicitation with neurological patients. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp. 158–168). New York, NY: Oxford University Press.

Levenson, R. W., & Ekman, P. (2002). Difficulty does not account for emotion-specific heart rate changes in the directed facial action task. *Psychophysiology*, *39*, 397–405.

Levenson, R. W., & Gottman, J. M. (1983). Marital interaction: Physiological linkage and affective exchange. *Journal of Personality and Social Psychology*, *45*, 587–597.

Littleford, L. N., Wright, M., & Sayoc-Parial, M. (2005). White students' intergroup anxiety during same-race and interracial interactions: A multimethod approach. *Basic and Applied Social Psychology*, *27*(1), 85–94.

Lovell, D. (Producer), & Zeffirelli, F. (Director). (1979). *The champ* [Motion picture]. Beverly Hills, CA: MGM/Pathe Home Video.

McKeown, G., Valstar, M., Cowie, R., Pantic, M. & Schroder, M. 2012. The SEMAINE database: Annotated multimodal records of emotionally coloured conversations between a person and a limited agent. *IEEE Transactions on Affective Computing*, *3*, 5–17.

Merikle, P. M., & Reingold, E. M. (1998). On demonstrating unconscious perception: Comment on Draine and Greenwald (1998). *Journal of Experimental Psychology: General*, *127*, 304–310.

Mikels, J. A., Fredrickson, B. L., Larkin, G. R., Lindberg, C. M., Maglio, S. M., & Reuter-Lorenz, P. A. (2005). Emotional category data on images from the International Affective Picture System. *Behavior Research Methods*, *37*(4), 626–630.

Monkaresi, H., Hussain, M. S., & Calvo, R. A. (2012). Classification of affects using head movement, skin color features, and physiological signals. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics* (pp. 2664–2669). New York: Institute of Electrical and Electronics Engineers.

Monteith, M. J., Ashburn-Nardo, L., Voils, C. I., & Czopp, A. M. (2002). Putting the brakes on prejudice: On the development and operation of cues for control. *Journal of Personality and Social Psychology*, *83*, 1029–1050.

Myers, D. & Tingley, D. (2011). *The influence of emotion on trust*. (Unpublished PhD thesis.) Princeton, NJ: Princeton University.

Öhman, A. (1986). Face the beast and fear the face: Animal and social fears as prototypes for evolutionary analyses of emotion. *Psychophysiology*, *23*, 123–145.

Öhman, A., & Soares, J. J. F. (1994). "Unconscious anxiety": Phobic responses to masked stimuli. *Journal of Abnormal Psychology*, *103*, 231–240.

Philippot, P., Chapelle, G., & Blairy, S. (2002). Respiratory feedback in the generation of emotion. *Cognition & Emotion*, *16*(5), 605–627.

Picard, R. (1997). *Affective computing*. Cambridge, MA: MIT Press.

Reiner, R. (Producer/Director), Scheinman, A., Stolt, J., & Nicolaides, S. (Producers). (1989). *When Harry met Sally* [Motion picture]. Hollywood, CA: New Line Home Video.

Reppeti, R. L., & Wood, J. (1997). Effects of daily stress at work on mothers' interactions with preschoolers. *Journal of Family Psychology*, *11*(1), 90–108.

Roberts, N. A., Tsai, J. L., & Coan, J. A. (2007). Emotion elicitation using dyadic interaction tasks. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp.106–123). New York, NY: Oxford University Press.

Rolls, E. T. (2007). Emotion elicitation by primary reinforcers and following stimulus-reinforcement association learning. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp. 137–157). New York, NY: Oxford University Press.

Rottenberg, J., Ray, R. D., & Gross, J. J. (2007). Emotion elicitation using films. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp. 9–28). New York, NY: Oxford University Press.

Russell, J. (2003) Core affect and the psychological construction of emotion. *Psychological Review*, *110*, 145–172.

Ryan, L., & Eich, E. (2000). Mood dependence and implicit memory. In E. Tulving (Ed.), *Memory, consciousness, and the brain* (pp. 91–105). Philadelphia: Psychology Press.

Schaefer, A., Nils, F., Sanchez, X., & Philippot, P. (2010). Assessing the effectiveness of a large database of emotion-eliciting films: A new tool for emotion researchers. *Cognition and Emotion*, *24*(7), 1153–1172.

Schnall, S., Laird, J. D., Campbell, L., Swang, H., Silverman, S., & Sullivan, D. (2000). *More than meets the eye: Avoiding gaze makes you feel guilty*. Paper presented at the annual meeting of the Society for Personality and Social Psychology, Nashville, TN.

Shortt, J. W., & Gottman, J. M. (1997). Closeness in young adult sibling relationships: Affective and physiological processes. *Social Development*, *6*(2), 142–164.

Seibert, P. S., & Ellis, H. C. (1991). A convenient self-referencing mood induction procedure. *Bulletin of the Psychonomic Society*, *29*(2), 121–124.

Siegman, A. W., & Boyle, S. (1993). Voices of fear and anxiety and sadness and depression: The effects of speech rate and loudness on fear and anxiety and sadness and depression. *Journal of Abnormal Psychology*, *102*, 430–437.

Soleymani, M., Pantic, M. & Pun, T.(2012). Multi-modal emotion recognition in response to videos. *IEEE Transactions on Affective Computing*, *3*(2), 211–233.

Stevenson, R. A., Mikels, J. A., & James, T. W. (2007). Characterization of the affective norms for English words by discrete emotional categories. *Behavior Research Methods*, *39*(4), 1020–1024.

Strack, F., Martin, L. L., & Stepper, S. (1988). Inhibiting and facilitating conditions of facial expressions: A non-obtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology*, *54*, 768–776.

Tsai, J. L., & Levenson, R. W. (1997). Cultural influences on emotional responding: Chinese American and European American dating couples during interpersonal conflict. *Journal of Cross-Cultural Psychology*, *28*(5), 600–625.

Vastfjall, D. (2001/2002). Emotion induction through music: A review of the musical mood induction procedure. *Musicae Scientiae* (Special Issue), 173–211.

Velten, E. (1968). A laboratory task for induction of mood states. *Behavior Research and Therapy*, *6*(4), 473–482.

Vuoskoski, J. K., Eerola, T. (2012). Can sad music really make you sad? Indirect measures of affective states induced by music and autobiographical memories. *Psychology of Aesthetics, Creativity, and the Arts*, *6*(3), 204.

Vural, E., Cetin, M., Ercil, A., Littlewort, G., Bartlett, M., & Movellan, J. (2007). Drowsy driver detection through facial movement analysis. *Human–Computer Interaction*, 6–18.

Watson, D. (2000). *Mood and temperament*. New York, NY: Guilford Press.

Weis, R., & Lovejoy, M. C. (2002). Information processing in everyday life: Emotion-congruent bias in mothers' reports of parent-child interactions. *Journal of Personality and Social Psychology*, *83*(1), 216–230.

Westermann, R., Spies, K., Stahl, G., & Hesse, F. W. (1996). Relative effectiveness and validity of mood induction procedures: A meta-analysis. *European Journal of Social Psychology*, *26*(4), 557–580.

Wiens, S., & Öhman, A. (2007). Probing unconscious emotional processes: On becoming a successful masketeer. In J. A. Coan & J. J. B. Allen (Eds.), *Handbook of emotion elicitation and assessment* (pp. 65–90). New York, NY: Oxford University Press.

Williams, G. P., & Kleinke, C. L. (1993). Effect of mutual gaze and touch on attraction, mood and cardiovascular reactivity. *Journal of Research in Personality*, *27*, 170–183.

Winkielman, P., Berridge, K. C., & Wilbarger, J. L. (2005). Unconscious affective reactions to masked happy versus angry faces influence consumption behavior and judgments of value. *Personality and Social Psychology Bulletin*, *31*, 121–135.

Wöllmer, M., Kaiser, M., Eyben, F., & Schuller, B. (2012). LSTM-modeling of continuous emotions in an audiovisual affect recognition framework. *Image and Vision Computing*, *31*(2), 153–163.

Zajonc, R. B., Murphy, S. T., & Inglehart, M. (1989). Feeling and facial efference: Implications of the vascular theory of emotion. *Psychological Review*, *96*, 395–416.

Zara, A., Maffiolo, V., Martin, J., & Devillers, L. (2007). Collection and annotation of a corpus of human-human multimodal interactions: Emotion and others anthropomorphic characteristics. *Affective Computing and Intelligent Interaction*, 464–475.