Psychology of Language and Communication 2020, Vol. 24, No. 1

sciendo

DOI: 10.2478/plc-2020-0001

STEPHANIE HUETTE
University of Memphis, Memphis, United States

NEGATION POLARIZES AGREEMENT DYNAMICS DURING SENTENCE COMPREHENSION

In a forced-choice mouse-tracking paradigm, true and false statements (ranging from very true, to ambiguous, to very false) were tested in both affirmative and negated forms. Replicating prior research, mouse trajectories reveal subtle differences in a continuum of true to false statements. However, negation modifies this process, particularly for very true statements (i.e. Bread is not made from sand). The mouse trajectories were more curved with negated sentences, with end-points of the continuum of truth (very true and very false statements) having the greatest area under the curve. The proposed explanation is the pragmatic meaning of a negated statement such as *"Gummie bears are not alive"* is infelicitous, whereas a true statement *"People live on Earth"* is felicitous. This study reveals the online dynamics of processing these statements and possible confusion, particularly when very true statements contain a negation.

Keywords: Embodiment, Eye-tracking, Language processing, Metaphor & Figurative, Language, Pragmaticss

The mind can rapidly process language in many kinds of forms, but the word *not* seems to delay us at times. The truth of a statement can be evaluated as a sentence unfolds, but when negation is introduced without appropriate context, it slows down this process (Dale & Duran, 2011; Fischler, Childers, Achariyapaopan, & Perry, 1985; Hagoort, Hald, Bastiaansen, & Petersson, 2004; Nieuwland & Kuperberg, 2008). How negation is integrated and the time-course of processing remain questions that speak to how we comprehend language, integrate contextual information, and the role of pragmatics as a constraint on sentence processing.

Address for correspondence: Stephanie Huette, Psychology, University of Memphis, 369 Psychology Bldg, 38152, Memphis, United States. E-mail: shuette@memphis.edu

The meaning of a negated sentence is very clear in many cases, but highly pragmatically constrained. For example, in the sentence "*The door is not open*," the meaning is its most probable alternative: a closed door. This relies on prior knowledge about antonym relationships. Other cases are not quite so clear, as in "*The eagle is not in the sky*." These are called *multiple alternative negations* and can be interpreted in several ways. For example, the eagle could be nesting or perched on top of a fence. Adding further complications to this problem, we can insert negation into almost any sentence, including those that are not fully true or fully false, as in "*My laundry might not be dry yet*" or "*Humans are not logical*." In natural conversation, we use negation frequently in different contexts, and at times negation can leave the linguistic code underspecified and must incorporate context in the form of pragmatic or knowledge constraints in order to be comprehended.

In traditional linguistic terms, negation has two routes of meaning: semantic and pragmatic. For example, the phrase "*Air is not blue*," relies on previous knowledge that air is colorless in order to know that this is a true statement. If we choose to say, "*Air is not blue*" to another person, the pragmatic meaning may be that the listener has a belief that air is blue and that we are attempting to refute this. Additionally, the semantics of this sentence are underspecified, such that *not blue* has many possible alternatives. When something is simply *not blue*, the language itself does nothing to specify the logical alternative, whether it is red or green or colorless. Knowledge that alternatives to blue are other colors comes from having experience with the world, and we must use this prior knowledge, or concurrent contextual cues such as pointing or discourse context, to disambiguate the meaning. At a higher level of characterization, we might say that different types of context interact with and constrain the meaning of negated phrases. Prior work has addressed this complexity in negation in depth, but the dynamics during the processing of negation have yet to be explored in the context of a wide range of propositions ranging from true to false.

Some have hypothesized that negation must be understood in terms of its affirmative counterpart, called a *counterfactual* (i.e., Ferguson, Sanford, & Leuthold, 2008; Nieuwland & Martin, 2012). Counterfactuals are, generally, a set of events that could have been or might be. In the case of negation, it is simply the affirmative proposition. Those who ascribe to investigating linguistic negation in terms of counterfactuals tend to posit that in order to build a comprehension model of a sentence such as "*The eagle is not in the sky*," there must first be a comprehension of the counterfactual "*The eagle is in the sky*." However, from a pragmatic viewpoint, the comprehension can come from other sources of information like vision, the identity of the speaker, and the current discourse. Further, the counterfactual itself might be insufficient for comprehending the negated version. For example, "*John is not in his office*" could imply more than absence, perhaps signaling that John should be in his

office. Comprehension of the affirmative "*John is in his office*" would not be enough to comprehend this implied meaning (referred to as conversational implicature). As such, implicature may be modifying the meaning in negated sentences and altering perceptions of the meaning.

There are many questions that arise from the study of language itself and meaning outside the linguistic code. Here, I examine the question of *how* negation is processed and what that means for the principles of processing negated implicature. In negation, the implicature is often an alternative, absence, or prohibition. Further, sentences that are infelicitous (sound odd) in the negated form sound felicitous in their affirmative form, and vice versa. This asymmetry in perception is relatively unexplored, and the purpose of the following study is to investigate the dynamics of agreeing or disagreeing with a wide variety of sentences in both, affirmative and negated forms.

Motivation & Background

In the work by Huette (2016), mouse-tracking was used to explore affirmative and negated sentence dynamics. She found that negated sentences were processed more slowly, and trajectories curved *less* toward the competing representation. This evidence was contrary to theories that posit an affirmative-first or competing affirmative representation. In a computational model, it was shown that an accumulative model fit the results better than a competitive model. Negation took more time and trajectories were less curved than an affirmative sentence when in a felicitous context. According to Huette, this is because when pragmatically constrained (two-alternative pictures to match a sentence with), it takes more time to consider the full context, even though both potential responses appear to be considered in parallel. That is to say, the conclusion was that the internal lexical-semantic representations were not competing, but rather the context was being incorporated on the fly to assemble the meaning of "*The newspaper is not on the rack*" with a picture of a newspaper on a rack and another of it in a driveway, or "*The towel is not on the floor*," where there is a towel on the floor and a towel on a towel bar.

Dale and Duran (2011) investigated negation using computer mouse-tracking for affirmative and negated statements such as "*Elephants are (not) large*" and "*Elephants are (not) small*." In this study, they found an increase in the number of times participants switched direction along the x-axis of the computer screen. Acceleration and deceleration events were also increased on trials with negated sentences. They also found that when pragmatically licensed by a preceding context, the increase in x-flips and acceleration and deceleration events went away. They did this by including a context of an adult speaking to a child and saying, "You want to lift an elephant?" followed by "But elephants are not small." Sufficient context facilitates processing of

negated sentences. This kind of modulation by context has been shown for negation in other ways as well (i.e., Nieuwland & Kuperberg, 2008).

Computer mouse-tracking has been used in a number of studies, ranging from syntactic parsing to bilingualism (Farmer, Cargill, Hindy, Dale & Spivey, 2007; Incera & McLennan, 2016). In a similar vein as negation, Duran, Dale and McNamara (2010) investigated the response dynamics of deception using a Nintendo Wii remote. They found that deceptive *yes* responses carried more complex movement signatures than truthful responses. The trajectory of a deceptive *yes* also veered closer to the truthful competitor than deceptive no responses, indicating an asymmetry between *yes* and *no* responses.

While there are many kinds of metrics such as velocity, maximum deviation, and x-flips that can be derived from a mouse trajectory, the area under the curve has been shown to provide a window into decision-making processes that are continually unfolding as they happen (Magnuson, 2005). Maldonado, Dunbar and Chemla (2019) found that spatial information was more informative than temporal information in Dale and Duran's (2011) negation study. Spatial deviations were also found to reflect processing of scalar implicature, where discrete shifts in mouse trajectories were found (Tomlinson Jr., Bailey, & Bott, 2013). Whether the trajectories are smoothly unfolding or discretely switching reflects the underlying processing: probabilistic or discrete. Both have been found in the mouse-tracking literature and thus, either could be found in the processing of a continuum of truth.

Other previous research has demonstrated that negated sentences prime an affirmative depiction of a sentence in the absence of any alternatives (Kaup, Yaxley, Madden, Zwaan, & Lüdtke, 2006). In this task, participants were presented with a sentence and then asked to respond to a picture. What the researchers found was that after reading "*The eagle is in the sky*," participants were faster to respond to a picture of an eagle in the sky, and slower to a picture of an eagle perched in a tree. When the sentence was "*The eagle is not in the sky*," participants were, again, faster to respond to a picture of an eagle in the sky. Participants were presented with the pictures in isolation, and visual context is well-known to be rapidly and immediately incorporated into sentence processing (Tanenhaus, Spivey-Knowlton, Eberhard & Sedivy, 1995). However, because only one picture was presented in isolation, and the task was to respond whether the object of the sentence was mentioned at all (not necessarily matching the sentence location), the experiment fails to show what happens in a more pragmatically-guided situation where alternatives are present. The experiment likely induced a "*Don't think of pink elephants*" kind of thinking because there were no alternatives available, so "*The eagle is not in the sky*" primed a picture of an eagle in the sky. When alternatives are present, a different kind of processing may be occurring.

For example, the representation of "*The towel is not on the floor*" is a towel on the floor, unless given alternatives.

In everyday use of language, we often agree and disagree with others, particularly when they make statements that can be evaluated based on the semantic meaning, such as "*Wolves are carnivores*." If a person has prior knowledge about wolves, she may be able to agree with this statement immediately. If a person has no prior knowledge about wolves, she may comply immediately with Grice's maxim of quality (the assumption that speakers will, generally, speak truthfully), and also agree with this statement. However, if a listener has knowledge contrary to this statement, regardless of its accuracy, she may disagree outright or simply hesitate to accept this statement. When negation is used, these conditions change. The phrase "*Wolves are not carnivores*" becomes a false statement, and hence, the truth-value of the statement has been inverted. However, I argue here, with empirical evidence, that this inversion of truth does not happen perfectly or symmetrically because of pragmatic constraints that affect the processing of negation.

Stating something in the affirmative or negated form, even if logically equivalent, may carry differences in meaning. For example, "*The door is not open*" is logically equivalent to "*The door is closed*" but the negated sentence may be a pragmatically-driven choice that has implications for comprehension. Here, we can look at existing work on so-called *truthiness*, which is defined as sentences containing ambiguity or partial truth as rated by human participants. McKinstry, Dale and Spivey (2008) demonstrated probabilistic responding to sentences that indicated the perceived truth value. Using a mouse-tracking paradigm in which participants rated each sentence as "True" or "False," participants' response trajectories for very untrue statements (e.g., "*One thousand is more than one million*") stayed closest to the "False" response, while ambiguous statements (e.g., "*Murder is sometimes justifiable*") averaged along the midline of the screen, between the "True" and "False" response areas.

Instead of focusing on the perceived truth or falsity of such statements, this work aims to gauge agreement. Because agreement is more subjective than "true" or "false," this will tap into processes of everyday conversational dynamics better than truthfulness, which is a more objective, logical judgment. Assent that a phrase is understood or agreed with comes in many forms in conversation, including head nods and backchannel communication (i.e., "Yeah" or "Uh huh") which are both forms of agreement (Dittman & Llewellyn, 1968). Thus, agreement and disagreement are the focus of this study in order to mimic conversational conditions more closely.

In order to understand the underlying processes of negation and to what degree it affects agreement and disagreement dynamics, I needed to investigate a similar continuum of stimuli. McKinstry et al. (2008) demonstrated that truth is not processed symbolically in terms of true or false,

but rather gradations in truth are maintained and used during processing. Gradient sensitivity is present in many types of categorization and decision processes (Dale, Kehoe, & Spivey, 2007; Huette & McMurray, 2010; McMurray, Tanenhaus, & Aslin, 2002). Under what I call here the *polarized gradient negation* hypothesis, the endpoint response behavior can appear to be a discrete agreement or disagreement as it must be by virtue of having two response options. However, during processing, the agreement or disagreement with a negated sentence should exhibit a steeper *gradient* (or slope) for relationship between an x-axis describing the stimulus truth value and a y-axis that includes a behavioral index of agreement.

The stimuli used in the current work were sentences that normed in terms of their truth. The statement "*Wolves are carnivores*" was normed to be 0.8 true, so one would assume that the inversion of this would make the negated statement 0.2 true on a probabilistic scale. However, the pragmatic use of negation may distort the meaning and thus how much it is agreed with. "*People live on Earth*" is a very true statement. Thus, it is extremely likely that one will be able to readily disagree with the statement "*People do not live on Earth*." The same should hold for the opposite, very true statements, such as "*Bread is not made from sand*." Logically, the negated is true and the affirmative version is false, but the negated contains a pragmatically odd presupposition. What the *polarized gradient negation hypothesis* predicts is the pragmatically odd negated statements, especially the very true statements, will have more curvature than the comparable affirmative statements. The alternative is that participants will treat both as logical statements, and treat both affirmative and negated statements equivalently, resulting in the same pattern of mouse movements for affirmative and negated statements.

If negated statements are agreed with more, this would be a compensatory process where the oddness or infelicity of the presupposition caused one to over-compensate for how true it is. For example, "*Bread is not made from sand*" presupposes that it is possible someone believes that bread can be made from sand, and the listener's response is to disagree to a greater degree than they would for an affirmative sentence of the same truth value. Greater curvature in trajectories would result in more area under the curve, particularly for the endpoints of the continuum, which would manifest as the gradient being steeper for negated sentences. If negated statements are agreed with less, this would mean the affirmative meaning is also being activated and being considered. In this case, "*Bread is made from sand*" would require activating the concept of bread made of sand first, representing the affirmative before a negation comes in to represent it. This is would result in the negated sentences exhibiting a shallower gradient.

These differences are predicted to be in how different the truth-continuum steps are from one another, with negation producing greater differences between steps (measured by the slope of the area under the curve by

continuum step). A steeper slope would indicate greater differences between the steps while a shallower slope indicates less differences in the curvature between steps. A steeper slope would indicate an expansion of the truth-space, where statements that are negated resulting in a very true statement (e.g., "*Bread is not made from sand*") are more obviously and strongly true than their affirmative counterparts (e.g., "*People live on Earth*"). This would provide support for the polarization of negation. A similar pattern of results should emerge for negated statements that are false: "*People do not live on Earth*" should be processed as more disagreeable than "*Bread is made from sand*." However, if the results show a shallower slope, this would indicate that negated statements are perceived in terms of both their counterfactual and the factual, supporting the idea that counterfactuals are evoked during negation processing and this results in less polarized agreement or disagreement.

Method

Participants

Participants volunteered for course credit or extra credit in an undergraduate course at the University of California, Merced. Eighty-four participants were run in a quiet room, and voluntary consent was obtained. All participants were native English speakers, right-handed, and had neither hearing nor (uncorrected) visual impairments, as indicated by self-report. Sample size was determined by Sampling Size of Population Mean, estimating the total number of eligible participants in the participant pool with a confidence level of 95% and a margin of error of 10%, yielding a sample size of 84 participants out of the estimated eligible 675 participants. The study was approved by the Internal Review Board of the University of California, Merced.

Materials

Six sentences at each truth-value were chosen from the Mindpixel database used in McKinstry et al., (2008) on the criteria that they were grammatically correct and easily negated in a similar place in the sentence. This database is a large collection of user-generated sentences that were then rated as either true or false by other users. This database is available from the first author upon request. The stimuli chosen were different from those used in McKinstry et al., using a finer-grained 11-point continuum as well as six stimuli at each continuum step instead of one. Each sentence was recorded by a male speaker with a Midwestern accent at the University of Cincinnati. Each stimulus was recorded in both affirmative and negated forms, for a total of 132 stimuli. All stimuli were recorded in a monotone voice and at an even pace, so as not to emphasize the negation in any way. Careful inspection of the stimuli was done to ensure sentences were as equivalent as possible in

both affirmative and negated forms. A transcription of all stimuli is included in the Appendix. PsyScope was used for data collection (Cohen, MacWhinney, Flatt, & Provost, 1993).

Design

Lists were created such that half of the stimuli were presented to half of the participants in the affirmative version, and that same half was presented in the negated form to the other half of participants. This ensures that the same sentence in both affirmative and negated forms is not repeated to one person. Each participant responded to 33 affirmative and 33 negated sentences in random order. This is a total of 66 sentences, or half of the total 132 sentences, with all 132 being used across participants. Three of the sentences were Affirmative and three were Negated at each continuum step. The two stimuli lists were crossed with the side of the screen, such that each stimulus list had half of participants with "Agree" on the left and the other half with "Agree" on the right side of the screen.

Procedure

Each participant was seated at a computer with a computer mouse and plenty of room for a full range of arm-motion. The tracking speed of the computer was set to the second lowest gain setting and acceleration was turned off, so the cursor movement on screen was approximately the same amplitude as the physical motion. This prevents ballistic responses and noise in the data (e.g., moving the mouse up and down vertically on the screen while thinking). The participant was fitted with over-ear headphones and given both oral and written instructions to click a button at the bottom center of the screen before moving upward toward a response. Participants were not given any information about how to move the cursor. After participants clicked the button, the stimulus played over the headphones and response options "Agree" and "Disagree" appeared. Participants were required to click the response to move on to the next trial. The computer recorded *x*,*y* pixel coordinates at 60 Hz as well as response, stimulus heard, and movement durations.

Replicated Analyses and Results

The following analyses look at response patterns as well as a measure of the trajectories and movement durations. In the following analyses, all data is included; no participants or trials were excluded for any reason. Responses were not classified as incorrect or correct due to the ambiguous nature of many of the trials. Additionally, I avoided discarding trials because participants likely responded very differently, meaning that too much data would be lost in the process. All data can be found at: https://osf.io/6kfvd/.

The mean proportion of Agree responses when the sentence was heard in the Affirmative was 0.54 (SD = .08) and 0.46 when the sentence was heard in

the Negated form (*SD* = 0.09). A pairwise comparison of means revealed this to be a significant difference: Overall, participants responded more with "Agree" when sentences were presented in the Affirmative, $t(83) = 5.5$, $p < .0001$. This replicates previous results that demonstrate a bias toward responding positively (see McKinstry et al., 2008). Participants tended to exhibit a greater yes-bias when stimuli were presented in the Affirmative sentence form. Movement durations also replicated previous results, demonstrating that Negated sentences (*M* = 4665ms, *SD* = 1039ms) take longer to respond to than Affirmative (*M* = 3917ms, *SD* = 808ms). A paired t-test showed a significant difference between the two, $t(83) = 14$, $p < .0001$.

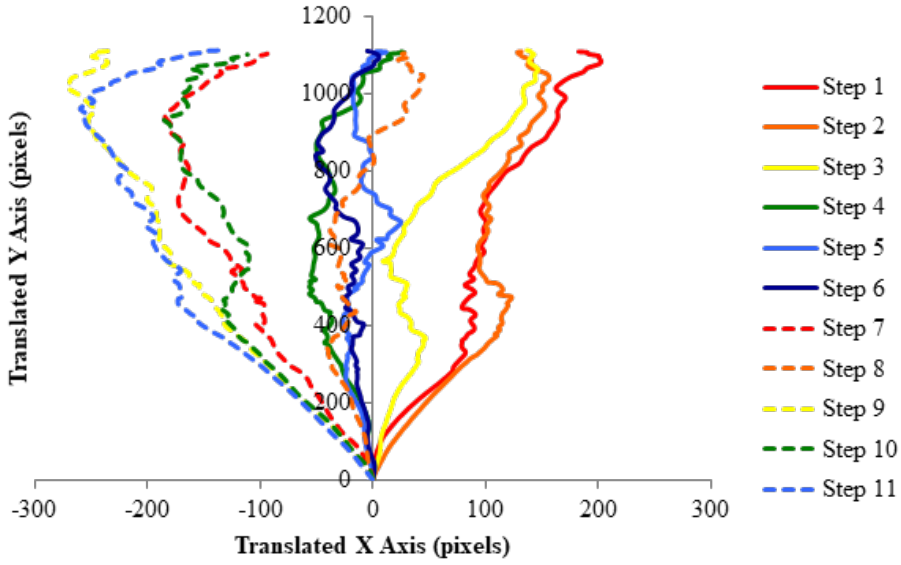

Figure 1. Averages of mouse trajectories, Affirmative sentences: The upper right was the Disagree target, and the upper left was Agree. Very false statements (Step 1) deviated very little whereas more ambiguous stimuli averages are found near the center. Very true statements also deviated very little (Step 11).

Response trajectory coordinates were translated such that all trajectories began at the common coordinate (0,0), thus getting rid of variability in the initial position of the mouse. Next, the two conditions where the Agree response option was on the right side of the screen were mirrored and reversed, so that all responses reflected Agree on the upper left and Disagree on the upper right. In order to analyze data with a variable number of points, *x*,*y* coordinates were time-normalized to 101 time-steps using linear interpolation (Huette & McMurray, 2010; Spivey, Grosjean & Knoblich, 2005). Averages for the Affirmative stimuli can be seen in Figure 1 and for Negated stimuli in Figure 2. Figure 1 shows a line for each step, where Step 11 is a very true affirmative sentence, and Step 1 is very false (see Appendix for stimuli).

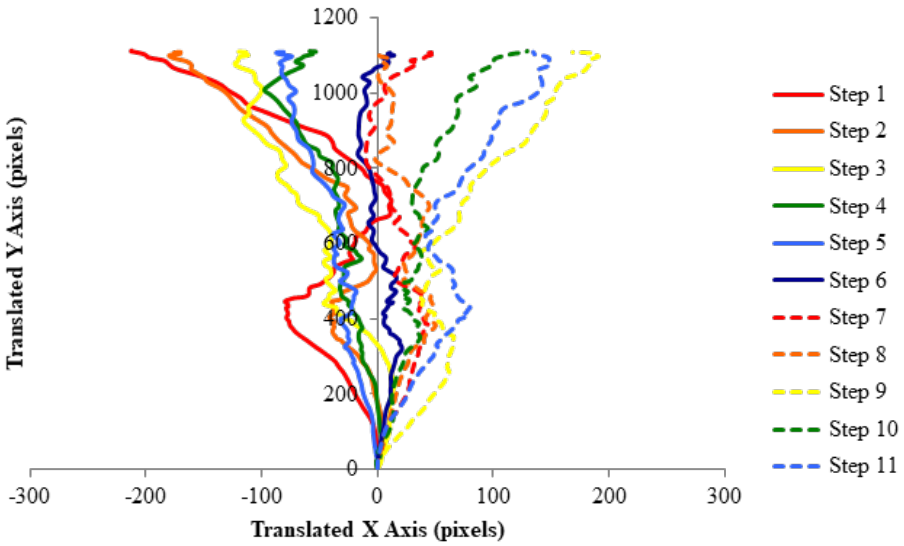

Figure 2. Averages of mouse trajectories, Negated sentences: Trajectories were averaged by continuum step. Note that the step values follow an approximate inverse pattern, where Step 1 of the continuum stays close to the upper left Agree area and Step 11 adheres closely to the Disagree area.

McKinstry et al. (2008) only analyzed final x-coordinate as seen in Figures 1 and 2. In order to more fully capture the entire trajectory, they were rotated 90-degrees clockwise, and the area under the curve (AUC) was computed. Importantly, before the dependent variable was computed, a fixed number was added to every x-coordinate so that all numbers were positive, avoiding averaging in negative numbers or zeros which might distort results. ANOVA was used as it was in McKinstry et al. in the preliminary analyses.

This area under the curve reflects the closeness to the Disagree response option, where lower values indicate a greater closeness to Disagree on average than higher numbers. Thus, Step 1 from the Affirmative continuum, rotated 90-degrees clockwise, would exhibit a lower average area under the curve than Steps 2, 3, 4, etc. As the area under the curve increases for the Affirmative condition, the closer the participants veered toward the Agree response option. Degree of Agreement and area under the curve was predicted to be expressed as a positive linear relationship. Contrarily, because negation inverts the continuum, as seen in Figure 1, Step 11 should exhibit the lowest area under the curve. Thus, Negated stimuli will show a negative linear relationship, demonstrating that participants veer closer and closer to Disagree as statements become more false. Again, the proposed hypothesis is reliant on this linear relationship, so first, the linearity was statistically tested with significance testing of trend using ANOVA. Secondly, the slope of the

Negated condition was proposed to be steeper, which was addressed by fitting a line to each participant's data, determining the slope, and using this as the dependent value.

To test whether the area under the curve increased linearly with each step, a mixed 11x2 ANOVA was used with Step (11) as a within-subject variable and Side of Response (2) as a between-subjects variable. This was done separately for Affirmative sentence trials and Negated sentence trials. The area under the curve increased linearly for Affirmative trials from Step 1 to Step 11 as predicted, linear trend: $F(1,82) = 201$, $p < .0001$. For negated trials, this trend was reversed, increasing linearly from Step 11 to Step 1, linear trend: $F(1,82) = 462$, $p < .0001$. Figure 3 depicts these results, along with equations for the slope and fits to a line. There was an effect of condition for both Affirmative, $F(1,82) = 7.7$, $p < .01$, and Negated trials, $F(1,82) = 6$, $p < .05$. The interaction was not significant, $F < 1$.

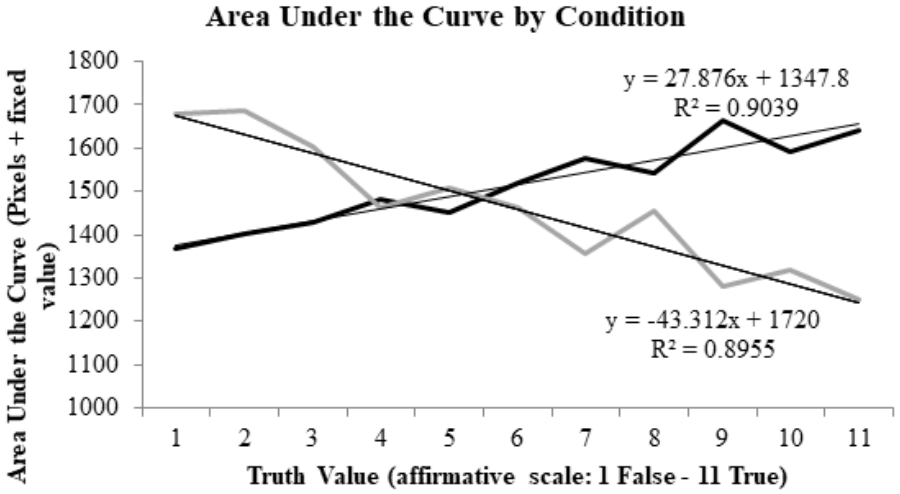

Figure 3. Area under the curve by condition: Significant linear trends are seen for both Affirmative sentences (black line) as well as Negated (grey line). Regression lines for both conditions are included. Negated statement step values are 1 True - 11 False.

To test differences in the polarization or shrinking of differences between steps in the different conditions, the slopes of each function were tested as the dependent measure. While both conditions exhibited a linear trend, differences in the steepness reflect differences in how agreeable or disagreeable these statements were perceived to be. A measure of how much the area under the curve changes from step to step was computed for each subject, for each condition (Affirmative and Negated) by computing a regression line using the method of least squares. Because the trends go in opposite directions (Affirmative: positive sloped; Negated: negative sloped),

the absolute value of the slope was used. This ensures that the statistical significance in difference between the two slopes is in the steepness, and not the directionality of the slope. A paired t-test showed the Negated condition (M = 43.7, SD = 17.6) was significantly steeper than the Affirmative (M = 28.1, SD = 17.8), $t(83) = -11.7$, $p < .0001$.

Extended Results

Because the response (Agree/Disagree) was included in the averages of the trajectories, further analysis was done to ensure this effect holds when broken down into 2x2 of Affirmative/Negated by Agree/Disagree. It could be that the response proportion was the driving force rather than curved trajectories in McKinstry et al. (2008), therefore, more in-depth analyses were warranted. The average trajectories are in Figures 4-7. Due to lack of data, only half of the response continuum is included in the averages and analyses (i.e., very few Agree responses to very false statements and vice versa).

In order to assess whether the trajectories differed significantly along the continuum of truth values, the area under the curve was calculated the same way and used as the dependent measurement. While ANOVA is straightforward and easy to interpret, it can be problematic for multiple reasons (Barr, Levy, Scheepers, & Tily, 2013; Brysbaert & Stevens, 2018). Linear mixed effects modeling was chosen as the alternative analysis for this data. I used the lme4 package for R to perform four linear mixed effects analyses of the relationship between truth value of the stimuli and the area under the curve (AUC) of responses (Bates, Maechler, & Bolker, 2012). AUC served as the outcome variable, Truth Value steps were entered as fixed effects, and participants were entered as random intercepts. Visual inspection of residual plots did not reveal any obvious deviations from homoscedasticity or normality. P-values were obtained by likelihood ratio tests of each model including the Truth Value effects against the corresponding Negation and Agreement model configuration without them.

In the Affirmative Agree condition, Truth Value significantly affected the AUC ($\chi^2(1) = 25.5$, $p < .001$), lowering it by about 12.7 ± 2.5 (standard errors) units per step. In the Affirmative Disagree condition, Truth Value significantly affected the AUC ($\chi^2(1) = 10.9$, $p < .001$), lowering it by about 10.2 ± 3.1 units per step. In the Negative Agree condition, Truth Value did not significantly affect the AUC ($\chi^2(1) = 0.1$, $p = .90$). In the Negative Disagree condition, Truth Value significantly affected the AUC ($\chi^2(1) = 4.7$, $p < .05$), raising it by about 6.5 ± 3.1 units per step. The direction of the effects matched our predictions in each case except for the Negative Agree condition, which did not yield a significant relationship.

Like the slope analysis done in the replication analyses, the slopes were analyzed but only for half of the continuum, as presented in Figures 4-7. Zero values were interpolated, and a line was fit to each individual for

Affirmative Agree, Affirmative Disagree, Negated Agree, and Negated Disagree statements. The absolute value of m in y = mx + b was used as the dependent measurement. The absolute value was used because the directionality of the slope was not the question, rather the steepness, and this avoids averaging positive and negative values which would misconstrue the mean. The mean slopes and standard deviations were: Affirmative Agree 23.8 (20.5), Affirmative Disagree 40.1 (36.2), Negated Agree 28.2 (23.3), and Negated Disagree 28.6 (30.7).

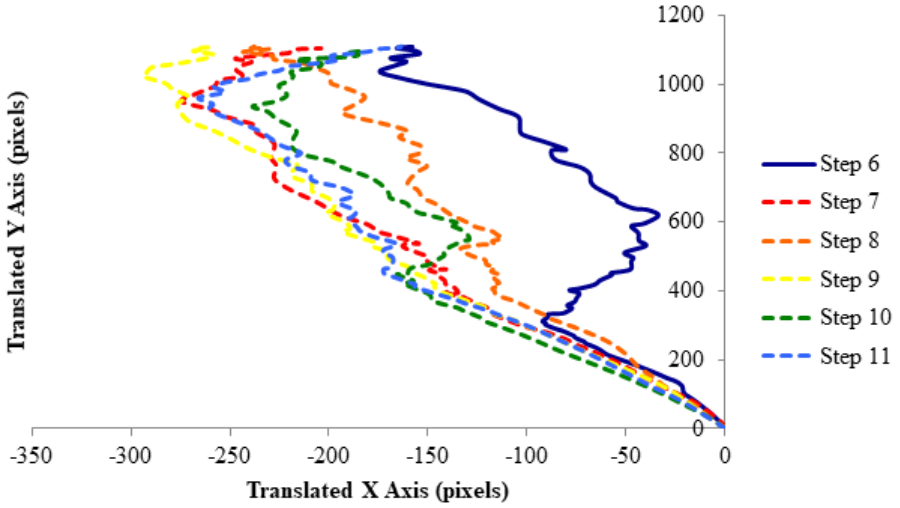

Figure 4. Affirmative sentences, Agree responses.

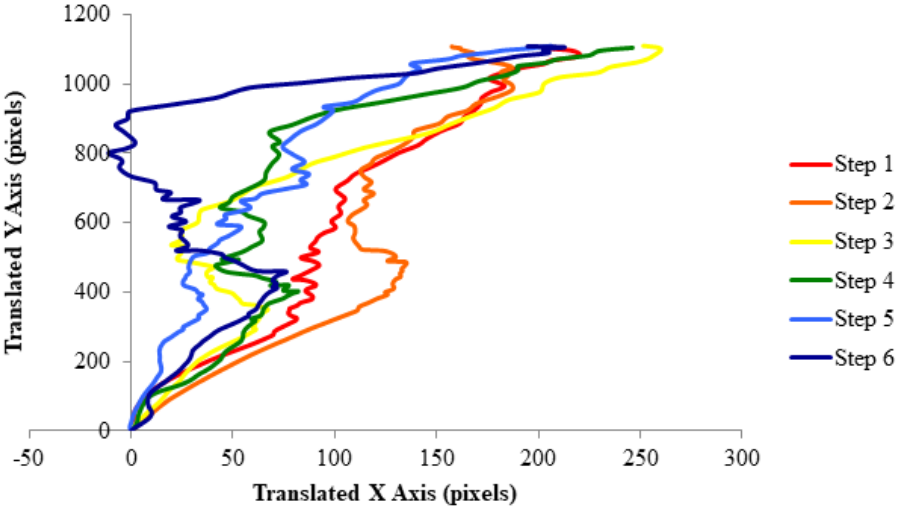

Figure 5. Affirmative sentences, Disagree responses.

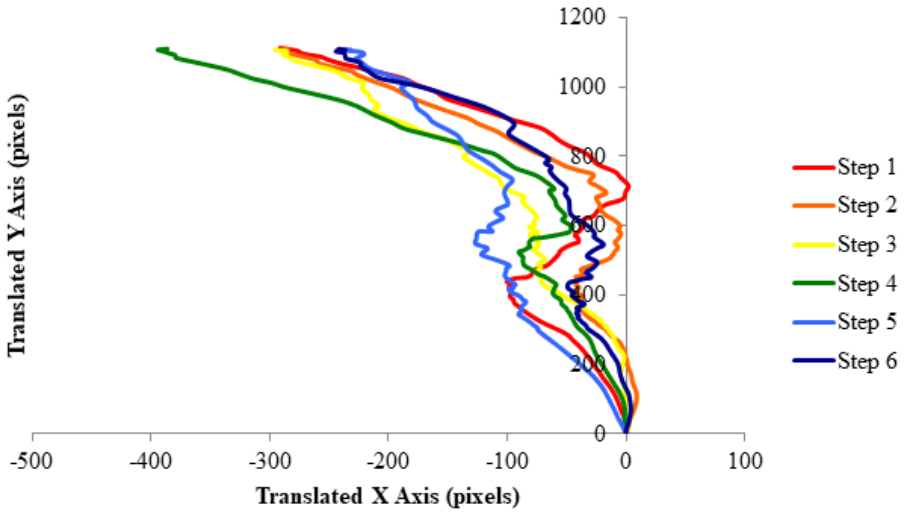

Figure 6. Negated sentences, Agree responses.

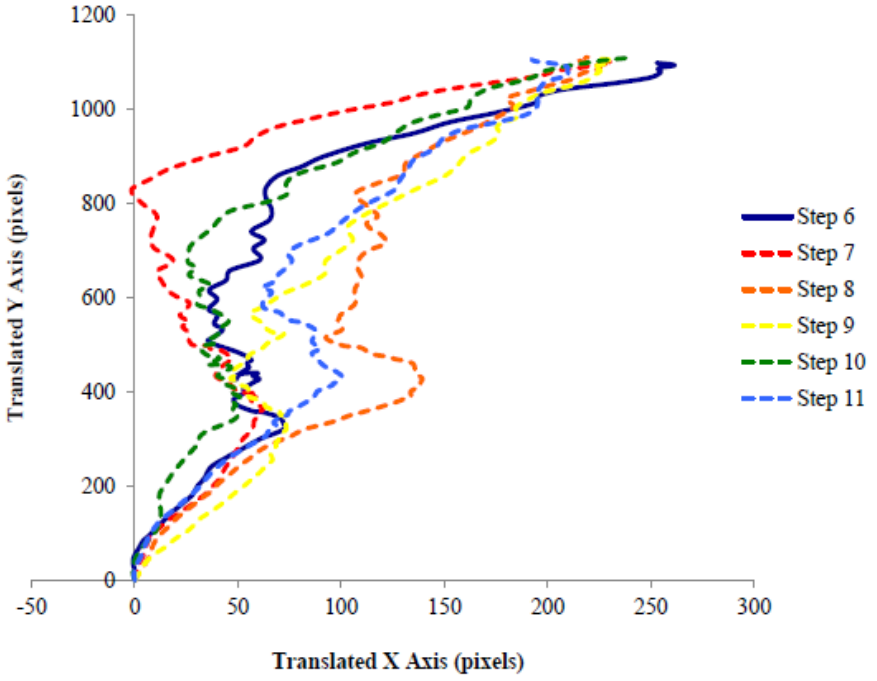

Figure 7. Negated sentences, Disagree responses.

In order to assess whether the Negation and Agreement conditions impacted the slope of participants' average trajectories, I constructed an additional linear mixed effects model on the data of all four condition combinations within the appropriate Truth Value steps. The outcome variable was the average slope of participants' mouse movement trajectories, while condition was entered as a fixed effect, and participants were entered as random intercepts. Again, inspection of residual plots did not reveal any violations of the assumptions of homoscedasticity or normality. I conducted a likelihood ratio test of this model against a null effects model without conditions entered as a fixed effect. Condition significantly affected the slope of participants' trajectories ($\chi^2(3) = 15.2$, $p < .01$).

Discussion

A continuum of stimuli varying by truthfulness was created and tested in both Affirmative and Negated forms. The mouse trajectories to responses demonstrate that both kinds of sentences preserve the degree of truthfulness of a statement but with different time courses and degrees of curvature in the trajectory. By showing this preservation of underlying truthfulness in the negated form, the current study's results provide evidence that support a gradient processing of both Affirmative and Negated linguistic forms. Both Affirmative and Negated sentences appear to have continuously unfolding dynamic trajectories, but manifest differently. Categorical phenomena such as phoneme recognition, color perception, and even types of animals carve a perceptual space far more discrete than the results seen here; however, these processes have also been shown to carry fine-grained details about the continuous physical dimension throughout the categorization process (e.g., Dahan & Tanenhaus, 2004; Dale et al., 2007; Huette & McMurray, 2010; McMurray & Aslin, 2005; McMurray et al., 2002).

The area under the curve was found to be linearly increasing or decreasing with truth value, demonstrating a gradient processing of truth values. However, when response was included in the analysis (Agree/Disagree), the Negated Agree statements such as "*Humans do not have sixteen legs*" were not significant across truth values. Although non-significant, this is a different trend from the Affirmative Agree statements. Therefore, the Negation may be doing something different to the truth value, particularly because almost all the stimuli were normed in an affirmative form.

What is relevant and informative for future research is how and in what ways mental dimensions becomes distorted or transformed through language. While it was demonstrated here that negation uniquely affects the speed and dynamics of agreement, the present research reveals a novel similarity in the underlying processes between affirmative and negated sentences, suggesting that there are not different mechanisms at work, but

that implied meaning is readily incorporated and used as a signal during communication. Just as affirmative sentences contain grades of "truthiness," so also does negation. Concepts even as seemingly symbolic and high-level as negation are not all-or-none in nature, even if the overt response appears to carve this perceptual space discretely.

Further, there is an oddness about negating a phrase at the endpoints of the continuum used here. These negated endpoints made the true statements very true and the false statements very false, and were on average higher or lower in terms of the area under the curve measurement than their affirmative counterpart. In particular, it is readily observable that the widest variance in responding was found in Step 1 along the negated statements: for example, hearing "*Apples are not edible*" and responding Disagree. This greater variability could be for several reasons, but it is an important consideration at a pragmatic level for future studies that examine sentences in both affirmative and negated forms. It is clear that "*Apples are edible*" and "*Apples are not edible*" are not perfectly inverted opposites. Thus, it appears that negation engages a process that relies more on pragmatic constraints than objective processing of factual and counterfactuals. Further, I have demonstrated the nature of this processing, where meaning becomes more polarized for a listener in the negated form. Theoretically, this means that negation need not activate an affirmative representation first, and pragmatic constraints are rapidly integrated into meaning comprehension during sentence processing. We might speculate that part of the intention of a speaker, when they choose a negated phrase over an affirmative phrasing, is to get people to process the message more deeply by emphasizing alternatives, absence or prohibition.

All these meanings have behavioral implications: alternatives being a signal to look around, absence being a signal of something being moved or taken away that was once there, and prohibition being the inhibition of an action (e.g., no smoking). All of these are more complex behaviorally than affirmative sentences and require further exploration in order to understand how low-level behaviors such as eye-movements might affect higher-order cognition such as the comprehension of alternatives negation. Future work on these issues will highlight missing theoretical pieces in the current understanding of language processing.

Acknowledgments

I would like to thank Michael Spivey, Teenie Matlock, Chris Kello, and Alexandra Paxton for their comments on the drafts of this manuscript. Thank you to John Hollander for assistance with the analyses.

References

Abrams, R. A., & Balota, D. A. (1991). Mental chronometry: Beyond reaction time. *Psychological Science*, *2*, 153–157.

Anderson, S., Huette, S., Matlock, T., & Spivey, M. (2010). On the temporal dynamics of negated perceptual simulations. In F. Parrill, V. Tobin, & M. Turner (Eds.), *Meaning, Form, & Body* (pp. 1–20). Stanford, CSLI.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, *68*(3), 255–278.

Bates, D., Maechler, M., & Bolker, B. (2012). *lme4: Linear mixed-effects models using S4 classes. R package version 0.999999-0*.

Brysbaert, M., & Stevens, M. (2018). Power analysis and effect size in mixed effects models: A tutorial. *Journal of Cognition*, *1*(1): 9, 1–20.

Cohen J. D., MacWhinney B., Flatt M., & Provost J. (1993). PsyScope: A new graphic interactive environment for designing psychology experiments. *Behavioral Research Methods, Instruments, and Computers*, *25*(2), 257–271.

Dahan, D., & Tanenhaus, M. K., (2004). Continuous mapping from sound to meaning in spoken-language comprehension: Immediate effects of verb-based thematic constraints. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 30*(2), 498–513.

Dale, R., & Duran, N. D. (2011). The cognitive dynamics of negated sentence verification. *Cognitive Science*, *35*, 983–996.

Dale, R., Kehoe, C., & Spivey, M. J. (2007). Graded motor responses in the time course of categorizing atypical exemplars. *Memory & Cognition*, *35*(1), 15–28.

Dittmann, A. T., & Llewellyn, L. G. (1968). Relationship between vocalizations and head nods as listener responses. *Journal of Personality and Social Psychology*, *9*(1), 79.

Duran, N. D., Dale, R., & McNamara, D. S. (2010). The action dynamics of overcoming the truth. *Psychonomic bulletin & review*, *17*(4), 486–491.

Farmer, T., Cargill, S., Hindy, N., Dale, R., & Spivey, M. (2007). Tracking the continuity of language comprehension: Computer mouse trajectories suggest parallel syntactic processing. *Cognitive Science*, *31*(5), 889–909.

Ferguson, H., Sanford, A. J., & Leuthold, H. (2008). Eye-movements and ERPs reveal the time-course of processing negation and remitting counterfactual worlds. *Brain Research*, *1236*, 113–125.

Fischler, I., Childers, D. G., Achariyapaopan, T., & Perry Jr., N. W. (1985). Brain potentials during sentence verification: Automatic aspects of comprehension. *Biological Psychology*, *21*(2), 83–105.

Gold, J., & Shadlen, M. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, *404*, 390–394.

Hagoort, P., Hald, L., Bastiaansen, M., & Petersson, K. M. (2004). Integration of word meaning and world knowledge in language comprehension. *Science*, *304*(5669), 438–441.

Huette, S. (2016). Putting context into context: Sources of context and a proposed mechanism for linguistic negation. Language, Cognition and *Neuroscience*, *31*(8), 1000–1014.

Huette, S., & Anderson, S. (2012). Negation without symbols: The importance of recurrence and context in linguistic negation. *Journal of Integrative Neuroscience*, *11*, 295–312.

Huette, S., & McMurray, B. (2010). Continuous dynamics of color categorization. *Psychonomic Bulletin & Review*, *17*(3), 348–354.

Incera, S., & McLennan, C. T. (2016). Mouse tracking reveals that bilinguals behave like experts. *Bilingualism: Language and Cognition*, *19*(3), 610–620.

Kaup, B., Yaxley, R. H., Madden, C. J., Zwaan, R. A., & Lüdtke, J. (2006). Experiential simulations of negated text information. The Quarterly *Journal of Experimental Psychology*, *60*, 976–990.

Kaup, B., Zwaan, R. A., & Lüdtke, J. (2007). The experiential view of language comprehension. How is negation represented? In F. Schmalhofer, & C. A. Perfetti (Eds.), *Higher Language Processes in the Brain* (pp. 255–289). Mahwah, NJ: Lawrence Erlbaum Associates.

Magnuson, J. S. (2005). Moving hand reveals dynamics of thought. *Proceedings of the National Academy of Sciences*, *102*(29), 9995–9996.

Maldonado, M., Dunbar, E., & Chemla, E. (2019). Mouse tracking as a window into decision making. *Behavior Research Methods*, *51*(3), 1085–1101.

McKinstry, C., Dale, R., & Spivey, M. J. (2008). Action dynamics reveal parallel competition in decision making. *Psychological Science*, *19*(1), 22–24.

McMurray, B., & Aslin, R. N. (2005). Infants are sensitive to within-category variation in speech perception. *Cognition*, *95*(2), B15–B26.

McMurray, B., Tanenhaus, M., & Aslin, R. (2002). Gradient effects of within-category phonetic variation on lexical access. *Cognition*, *86*, B33–B42.

Nieuwland, M. S., & Kuperberg, G. R. (2008). When the truth is not too hard to handle: An event-related potential study on the pragmatics of negation. *Psychological Science*, *19*(12), 1213–1218.

Nieuwland, M. S., & Martin, A. E. (2012). If the real world were irrelevant, so to speak: The role of propositional truth-value in counterfactual sentence comprehension. *Cognition*, *122*(1), 102–109.

Spivey, M. J., Grosjean, M., & Knoblich, G. (2005). Continuous attraction toward phonological competitors. *Proceedings of the National Academy of Sciences*, *102*(29), 10393–10398.

Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*(5217), 1632–1634.

Tomlinson Jr., J. M., Bailey, T. M., & Bott, L. (2013). Possibly all of that and then some: Scalar implicatures are understood in two steps. *Journal of Memory and Language*, *69*(1), 18–35.

Appendix

Step	Affirmative	Negated
Step 11	People live on Earth.	People do not live on Earth.
	People sometimes lie.	People do not sometimes lie.
	Computers need electricity.	Computers do not need electricity.
	Apples are edible.	Apples are not edible.
	Latin is a language.	Latin is not a language.
	Mosquito bites itch.	Mosquito bites do not itch.
Step 10	Airplanes can fly through thunderstorms.	Airplanes cannot fly through thunderstorms.
	Humans are mortal.	Humans are not mortal.
	Oak trees are tall.	Oak trees are not tall.
	Solar cells need light.	Solar cells do not need light.
	Pretzels are good snack food.	Pretzels are not good snack food.
	People see with their eyes.	People do not see with their eyes.
Step 9	Sandpaper is abrasive.	Sandpaper is not abrasive.
	Acting is a profession.	Acting is not a profession.
	Spinach is a leafy green vegetable.	Spinach is not a leafy green vegetable.
	Wolves are carnivores.	Wolves are not carnivores.
	Sweat tastes salty.	Sweat does not taste salty.
	Smoking causes cancer.	Smoking does not cause cancer.
Step 8	Black is the absence of color.	Black is not the absence of color.
	A platypus is a mammal.	A platypus is not a mammal.
	It may be possible to break the speed of light.	It may not be possible to break the speed of light.
	Chocolate is harmful to dogs.	Chocolate is not harmful to dogs.
	Borscht is made from beets.	Borscht is not made from beets.
	Humans can live in peace.	Humans cannot live in peace.
Step 7	Britney Spears is blonde.	Britney Spears is not blonde.
	Vindaloo is a kind of Indian food.	Vindaloo is not a kind of Indian food.
	Money is necessary for commerce.	Money is not necessary for commerce.
	A koala is a marsupial.	A koala is not a marsupial.
	Violet is a type of blue.	Violet is not a type of blue.
	The internet is addictive.	The internet is not addictive.
Step 6	Cows have an even number of toes.	Cows do not have an even number of toes.
	The soul is separate from the body.	The soul is not separate from the body.

Step	Affirmative	Negated
	Humans are logical.	Humans are not logical.
	Green is the complementary color for red.	Green is not the complemen-tary color for red.
	English is the most popular human language.	English is not the most popular human language.
	It is possible to stop thinking.	It is not possible to stop thinking.
Step 5	Microwaves cause cancer.	Microwaves do not cause cancer.
	Chopped dates are fat free.	Chopped dates are not fat free.
	A catalyst changes in a chemical reaction.	A catalyst does not change in a chemical reaction.
	Alcohol is good for your health.	Alcohol is not good for your health.
	Fish can swim backwards.	Fish cannot swim backwards.
	Lying is necessary.	Lying is not necessary.
Step 4	Morals and ethics are the same thing.	Morals and ethics are not the same thing.
	Time is the same as space.	Time is not the same as space.
	People can see in the dark.	People cannot see in the dark.
	Wind is an element.	Wind is not an element.
	Polar bears eat penguins.	Polar bears do not eat penguins.
	All glasses are made of glass.	All glasses are not made of glass.
Step 3	Pi is a rational number.	Pi is not a rational number.
	Floors are always above ceilings.	Floors are not always above ceilings.
	Poker is a board game.	Poker is not a board game.
	UFOs are real.	UFOs are not real.
	Burping is painful.	Burping is not painful.
	Ducks can earn PhDs.	Ducks cannot earn PhDs.
Step 2	The moon is hot.	The moon is not hot.
	It is legal to rob a bank.	It is not legal to rob a bank.
	Gummy bears are alive.	Gummy bears are not alive.
	Wood conducts electricity.	Wood does not conduct electricity.
	Air is blue.	Air is not blue.
	Bananas turn red when they ripen.	Bananas do not turn red when they ripen.
Step 1	A feather is heavier than a car.	A feather is not heavier than a car.
	Humans have sixteen legs.	Humans do not have sixteen legs.
	Unicorns are real.	Unicorns are not real.
	Jupiter is the closest planet to the sun.	Jupiter is not the closest planet to the sun.
	Television is made from cheese.	Television is not made from cheese.
	Bread is made from sand.	Bread is not made from sand.