Towards a personalized collaborator for Human-Robot creative ideation

Adapting conversational strategies to increase creativity

LUCIE GALLAND, Département d'informatique de l'ENS, ENS, CNRS, PSL University, France and Intelligent Systems, Delft University of Technology, The Netherlands
MILENE GUERREIRO GONCALVES, Faculty of Industrial Design, Delft University of Technology, The Netherlands
CATHARINE OERTEL, Intelligent Systems, Delft University of Technology, The Netherlands

Humans engage in creatively demanding tasks every day. However, thinking creatively while alone and without external input can be difficult. As other human collaborators are not always easily available, robots can serve as a scalable addition to the ideation process. In this work, we investigate how social robots can improve creativity in ideation tasks. Studies have shown that good mood and social robot partners can have a positive impact on creativity. We explore the possibility that a social robot, adapting its conversational strategies to user's mood, could enhance creativity. To this end, we develop a reinforcement learning algorithm that can learn a custom transition function to adapt to each particular user. We apply this algorithm in the case of an ideation task. A simulated user is developed to test the policy. We find that the reinforcement learning algorithm performed significantly better than the random use of these strategies. This algorithm will later be tested with human participants using a Furhat robot.

Additional Key Words and Phrases: Human-Robot interaction, ideation task, creativity, social robot, reinforcement learning

ACM Reference Format:
Lucie Galland, Milene Guerreiro Goncalves, and Catharine Oertel. 2022. Towards a personalized collaborator for Human-Robot creative ideation: Adapting conversational strategies to increase creativity. 1, 1 (May 2022), 6 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION AND RELATED WORKS

1.1 Motivation

There is a growing body of work that introduces robots as collaborators in creative tasks with children [2, 3] and adults [11, 12, 15, 17]. These tasks require participants to come up with as many creative ideas as possible in a certain context, such as designing a rock zen garden [15], finding creative ideas for water and energy conservation [12], or designing creative tangram shapes [11]. Creativity is classically defined as an idea that is novel and effective [22, 24, 27]. Novelty refers to solutions that are unique, original, or uncommon, while effectiveness refers to an idea that works well, is useful, and appropriate. Computational tools can help participants in such ideation tasks by providing stimuli like examples, suggestions, pictures, etc. The use of stimuli in these tasks helps develop creative ideas [9]. Furthermore, the automated and proactive provision of these stimuli is important for participants to find new and more creative ideas [8, 23]. As humans communicate better verbally, it would also be more intuitive to provide those stimuli pro actively through dialog [28]. Previous studies [12, 15] have shown that an interaction with a social robot also improves the impact of stimuli. However, social behaviors can have an impact on the user's emotion [18]. Studies have shown that people in positive moods generate more creative ideas than people in negative moods [4, 14]. Therefore, this study investigates the effect adaptation to mood through conversational strategies has on the creative process.

In this paper, we propose to design a robot collaborator for a brainstorming task capable of providing stimuli by verbally and socially interacting with the user while adapting to its emotions.

1.2 Related Works

1.2.1 Social robot and creativity. There is a large body of work investigating the possibility of using robots as collaborators in the design process. In [17], a robot arm collaborates with the participant on a design task in which the blocks must be placed in the most efficient way possible. The arm did not show any social behavior, but participants often attributed social meaning to the pragmatic behaviors of the robot, leading to deception when their feelings were not taken into account. Moreover, previous studies [2, 3, 11] showed the positive effect of a social robot on creativity. In [2], the participants who interacted with the social robot also had a stronger desire to work with the robot.

1.2.2 Verbally Interactive stimuli tools. In [28], the authors interviewed designers to identify the design requirements for the acceptance of an agent as a teammate. They found that such an agent should be able to understand and communicate via natural language, have human characteristics (identity and social cues), and be explainable and controllable (taking into account feedback). However, none of the above works offers the ability to verbally interact with the social robots collaborators. The use of human-robot dialog as interactive stimuli-providing tools has been explored in [15], where the authors provide stimuli through a PowerPoint presentation or through a robot to participants designing a Zen rock garden. They found that providing stimuli in an interactive verbal setting enhances the creativity of the participants. In [12], the authors found that participants consistently contributed more ideas and ideas of higher quality when they perceived their teamwork partner as a bot.

1.2.3 Mood and Creativity. The idea of relying on emotion to improve the creativity of participants has been explored in a few works. In [16] the authors have developed a computational model to include social and emotional context as a facilitator of the generation of group ideas. They found that encouraging the participant during a brainstorming task can positively affect their mood and creativity.

Authors' addresses: Lucie Galland, Département d'informatique de l'ENS, ENS, CNRS, PSL University, Paris, France and Intelligent Systems, Delft University of Technology, Delft, The Netherlands, lucie.galland@ens.psl.eu; Milene Guerreiro Goncalves, Faculty of Industrial Design, Delft University of Technology, Delft, The Netherlands, m.guerreirogoncalves@tudelft.nl; Catharine Oertel, Intelligent Systems, Delft University of Technology, Delft, The Netherlands, C.R.M.M.Oertel@tudelft.nl.

2022. XXXX-XXXX/2022/5-ART $15.00
https://doi.org/XXXXXXX.XXXXXXX

In [6], feedback on the ideas of the participants is provided at each turn. The computed generated feedback has an impact on the participant mood, and that using first negatively biased feedback followed by positively biased feedback significantly improves the originality of the created idea. Therefore, an agent capable of gratifying or criticizing user ideas at the right time seems to have a positive impact on creativity.

1.3 Contribution

The above works show that social robots can be good partners for creativity and that keeping the user in a good mood can be important for the creativity of its ideas. However, to our knowledge, no work has explored the possibility of designing a social robot collaborator that provides stimuli and adapts to user emotions during a creative task. Therefore, we propose designing a conversational agent capable of adapting its conversational strategies to the user's emotion to improve creativity. We propose to develop a model capable of choosing from a pool of conversational strategies according to the emotional context to answer our research questions :

- **RQ1**: Does a robot that provides conversational strategies improve creativity over a nonsocial robot ?
- **RQ2**: Can adapting these conversational strategies to the user's current emotions improve creativity compared to a random use of these strategies ?

To answer our research questions, we developed a conversational agent composed of a rule-based dialog manager, a creativity estimator, and a conversational strategy selector. At each speaking turn, a conversational strategy selected from (None, Encouragement[16], Praising[6], or Critic[6]) to maximize estimated creativity is inserted before the dialog act chosen by the rule-based dialog manager (e.g. "That last idea was really good, does it give you other ideas ?"). The agent is evaluated using the Alternative Uses Task (AUT) [26], where participants need to come up with creative uses for an object (e.g.: Find an alternative use for a rope by a badly intended person). The agent will act as a coach and help the user by asking some questions, suggesting themes, and providing feedback. In the following of this paper, Section 2 presents the method used to estimate the creativity of user ideas, Section 3 presents the mental state and decision policy of the agent, the model is evaluated in Section 4.

2 CREATIVITY ESTIMATOR

In this section, we develop an online approximate measure of creativity to use as a reward in our reinforcement learning algorithm.

2.1 Definition of creativity for the alternative use task

The Alternative Use Task (AUT) is usually scored in terms of Fluency (number of ideas), Originality (how original the ideas are), Flexibility (how varied the ideas are) and Elaboration (how developed the ideas are) [26]. To meet the definition of creativity (novelty and effectiveness)[22], we add the Effectiveness dimension (how appropriate and useful the ideas are) by adding a context to the classical AUT requirements (e.g. the use should fit a badly intended person). We develop a measure of creativity to give a turn-by-turn reward to our conversational agent, so we need to evaluate separated ideas. Therefore, we remove the Fluency and Flexibility dimensions for the turn-by-turn evaluation (these two dimensions should, however, be analyzed afterwards). As such, the measure of creativity of an idea will be a combination of :

- **Novelty** : How novel the idea is.
- **Elaboration** : How developed the idea is.
- **Effectiveness** : Along two axes :
  - **Appropriateness** : How well the idea fits the given item (eg box or rope).
  - **Usefulness** : How well the idea fits the given context (e.g. the box is used by a child or the rope is used by a badly intended person).

In this task, an idea must meet two requirements to be effective: It must be appropriate (the idea must be feasible with the given item) and useful (the idea must meet the required context). These four dimensions of our measure of creativity will allow the agent to provide appropriate feedback to the user (facing a short answer, the agent can ask the user to elaborate on its ideas). To calculate the reward of the reinforcement learning agent, an online measure of creativity is developed for our AUT task.

2.2 Method and Dataset

The work of [5] shows that creativity scores correlate with the distance from the item to the response in an embedding space. Therefore, embedding spaces can capture elements relevant to creativity. However, this method gives a global creativity score, whereas we want to be able to distinguish the score along the different dimensions of creativity. The following models are built based on the AUT response data set provided in the Supplementary Material of [5] where the participants completed two AUT trials with two different items : box and rope. Responses were scored for creativity quality using the subjective scoring method. This dataset has been annotated on Prolific using the subjective scoring method for appropriateness (how well the idea fits to the object rope or box) and usefulness (how well the requirements are respected: if the use of a box is useful for a baby or if the use of a rope is useful for a badly intended person). Two raters per item and measure rated each idea on a Likert scale from 1 to 5. After rescaling the answer and removing outliers (rope useful : 0%,rope appropriate : 6%,box useful : 3.7%,box appropriate:12%) acceptable Krippendorff's alphas were computed for each of this tasks ($\alpha_{\text{rope useful}} = 0.667$,$\alpha_{\text{rope appropriate}} = 0.603$,$\alpha_{\text{box useful}} = 0.703$, $\alpha_{\text{box appropriate}} = 0.71$). These annotations are then averaged to create an appropriateness and usefulness score.

2.3 Elaboration

We define the elaboration score as the number of significant words in the response. A significant word is a word that is not in the list of spacy's stop words for English [10], we add to this list the words make, build, and use, and the item name to adapt it to our task. This elaboration score is correlated with the creativity score rated by [5] (creativity = 0.24*elaboration + β, p<2e-16)

Task	Top1 accuracy	Mean square error	Top2	Top3
rope useful	0.68	0.25	0.87	0.96
rope appropriate	0.65	0.24	0.90	0.97
box useful	0.73	0.32	0.83	0.87
box appropriate	0.73	0.9	0.73	0.80

Table 1. Bert Classifier results

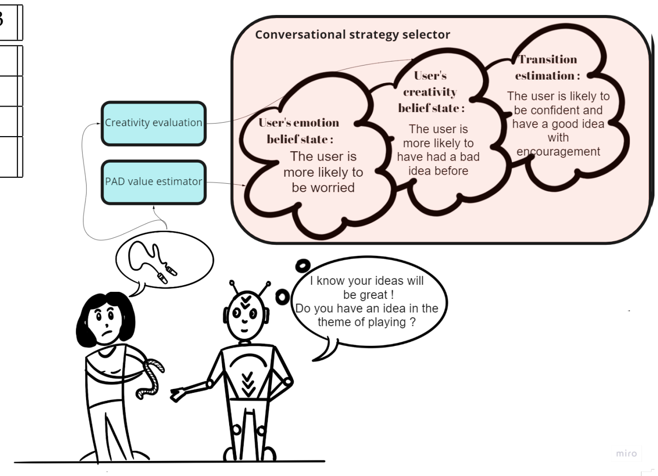

Fig. 1. Summary of the approach

2.4 Appropriateness and Usefulness

A classifier has been trained for each of these measures (usefulness and appropriateness) and each problem (rope and box) The classifiers is composed of a pretrained Bert model from the transformer library [29] followed by 3 linear layers interleaved with LeakyReLU layers. These classifiers reached an accuracy comparable to state-of-the-art classifiers for the sentence classification task. (see tab.1).

2.5 Novelty

Novelty is measured using an unsupervised approach inspired by code-based novelty measurement [1]. We use topic modeling to automatically apply code to the user response. We then attribute a score inversely proportional to the frequency of this code in the original dataset. Due to the short length of the responses, we use the gsdmm library [30] for short text classification. The linear regression between the creativity score (mainly based on novelty) and our unsupervised novelty score is significant with p<2e-16 and $\beta = 0.2$

2.6 Prediction of creativity score

Combining our elaboration, appropriateness, and novelty scores, we infer the creativity scores provided by [5] (the usefulness score is not taken into account since the context was not present in the study [5]). Each regression variable has a significant effect (p <2e-16) and $R^2 = 0.23$. These measures compute an approximation of the creativity of the answer that will act as a reward for our reinforcement learning algorithm.

3 AGENT'S MENTAL STATE

The agent's mental state is composed of a conversational strategy selector and a dialog manager that selects the next dialog act.

3.1 Conversational strategy selector

At each speaking turn, the agent selects a conversational strategy from (neutral, praise, encouragement, or critic). This choice will be based on the agent's perception of the current state, which includes the user's mood state and an estimation of the level of creativity of the user's idea, as well as the agent's estimation of the transition and reward functions associated with the word state (see Fig.1). Since moods are not directly observable and mood estimation can be noisy, we use a Partially Observable Markov Decision Process (POMDP). A POMDP is a tuple $\{S, A, T, R, \Omega, O\}$ where :

- $S = \{\text{mood}_{\text{hidden}}, \text{last used strategie}, \text{last idea quality}\}$ is the set of possible states. $\text{mood}_{\text{hidden}}$ can take 3 values : {Happy, Sad, Neutral}
- $A = \{\text{Neutral}, \text{Encourage}, \text{Praise}, \text{Critic}\}$ is the set of possible action
- $T : S \times A \times S \to [0; 1]$ is the transition function of the system ie. $p(s'|s, a) = T(s, a, s')$ the probability to be in state s' after taking action a in state s,it is unknown for the mood part of the state
- $R : S \times A \times S \to \mathbb{R}$ is the creativity score of the idea is the reward function and is known as idea quality is part of the state
- Ω_m : The Pleasure Arousal Dominance (PAD) space is the set of possible observations for mood
- Ω_c : The estimated creativity is the set of possible observations for creativity
- $O : S \times \Omega_m \times \Omega_c \to [0; 1]$ is the observation function that gives the probability of an observation in a given state $p(\omega_m, \omega_c|s) = O(s, \omega_m, \omega_c)$. It is determined using the OCC model that offers a mapping between PAD and the mood space [7] then $O(s) : \omega_m \times \omega_c \to softmax(-||s_{PAD} - \omega_m^{PAD}||) * (s_{creativity} - \omega_c)^2$ [13].

Throughout the interaction, the agent updates its belief state b, which is a probability distribution over all possible states: $b(s) = p(s|h)$ where h is the history.

3.1.1 Transition function. The transition function is unknown in our problem and can depend on the user: A user might enjoy an encouraging agent when the next user would not like it. A user might also have more ideas when relaxed, while another might thrive under pressure. This issue is addressed using α-POMDP[19] : the transition function is learned online and adapted to the user.

$$T_{\text{user}}(s,a,s') = \frac{L(s,a,s')}{N(s,a)}$$

$$N(s,a) = \text{number of time action } a \text{ has been taken is state } s$$

$$= \sum_{i=1}^{t} \mathbb{1}_{s_i=s, a_i=a, s_{i+1}=s'} b_i(s)$$

$$L(s,a,s') = \sum_{i=1}^{t} \mathbb{1}_{s_i=s, a_i=a} b_i(s')$$

$$T(s,a,s') = w_p T_{\text{user}}(s,a,s') + (1-w_p) T_{\text{mean}}(s,a,s')$$

And T_{mean} is updated at the end of every interaction:

$$T_{\text{mean}} = \frac{i_A * T_{\text{mean}} + T_{\text{user}}}{i_A + 1}$$

$$i_A = \text{Number of previous interactions}$$

3.1.2 Action selection. At each turn, we select the next action using the QMDP algorithm to optimize the reward $\alpha(R - 10(a_t == a_{t-1})$ where R is the reward for creativity. The next action is selected using a ϵ-greedy policy.

3.2 Dialog manager

The dialog manager selects the next dialog act of the agent that is independent of the conversational strategy used. A rule-based dialog manager is used to measure the effect of the conversational strategy selector.

4 EVALUATION

In the following evaluation, we simulate the interaction between a simulated user and our dialog agent with 3 different conditions.

- **Not social:** The agent follows the script and does not use any social strategies
- **Randomly social:** The agent follows the script and uses social strategies randomly
- **Adaptive social:** The agent follows the script and uses social strategies while adapting its policy to the particular user.
- **Adaptive social pretrained:** The agent follows the script and uses social strategies while adapting its policy to the particular user with a transition function initialized with the information obtained after 50 conversations under random conditions.

4.1 User simulator

The simulated user produces at each speaking turn a mood state and associated PAD values, as well as a creativity score and a corresponding natural language idea.

4.1.1 Simulated user's mood. The mood transition function of the simulated user is based on the OCC model of emotion that maps events to emotional reactions [25] (see fig 2). Each event causes a particular emotion in the user. Mood and PAD levels are inferred using the equations of [13].

4.1.2 Simulated user's idea. The frequency and quality of user ideas are determined based on the context of interaction and theory: According to [20] users in good mood are more creative and according to [6] critics can enhance creativity. The simulated user will thus have ideas of quality proportional to its mood and will have a better chance to have good ideas if a critic has been issued while the user is in a good mood. User ideas will be extracted from our database based on the quality of user ideas determined from its current mood and agent behavior.

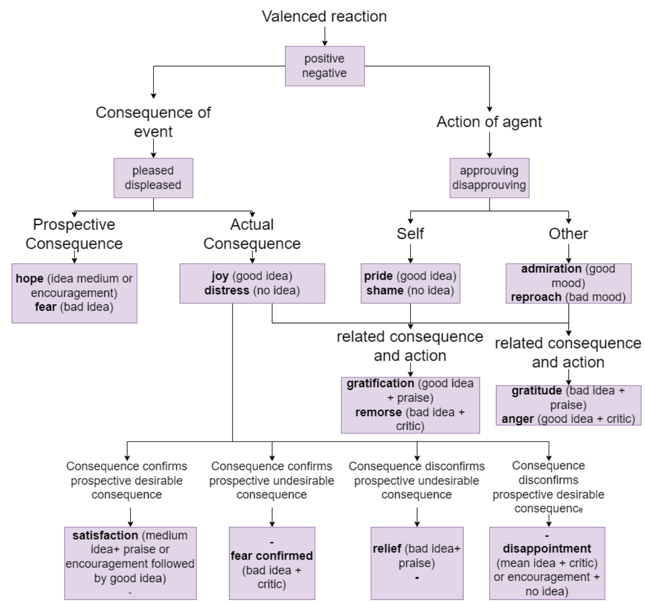

Fig. 2. Emotion reactions to events based on the OCC model[25]

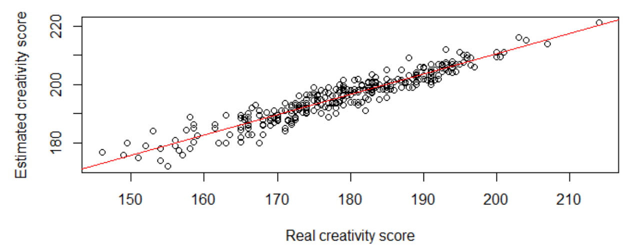

Fig. 3. Linear relationship between real and estimated idea quality

4.2 Results

The model was tested in interaction with 1 type of simulated user for 60 turns. The reward is computed using our creativity estimator (see Sect.2). There is a highly significant linear relationship between the quality of the real idea of the simulated user and the estimated one (p-value <2e-16, r = 0.9, see Fig.3). We found that our model is capable of adapting to the user simulator and significantly improving the quality of the simulated user ideas. The rewards are normally distributed (Shapiro test p-value = 0,25,0.78, 0.3 and 0.93), the variances are not equal between conditions (Bartlett test p-value

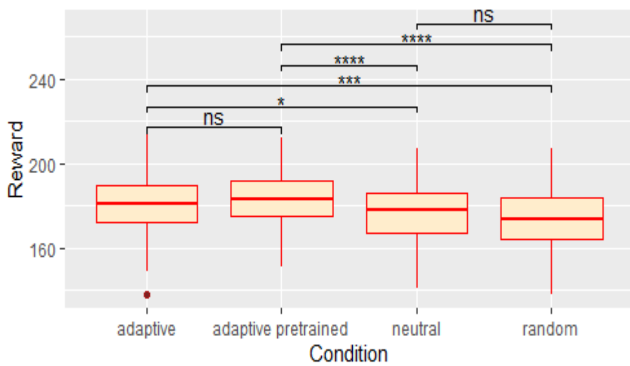

Fig. 4. Total reward per condition
ns : not significant,* p <0.1,** p < 0.01,*** p < 0.001,**** p < 0.0001

= 0.03) but the samples are of the same size for each condition. Therefore, we run a one-way analysis of variance (ANOVA) to compare the effect of the condition on the quality of the idea. A one-way ANOVA revealed that there was a statistically significant difference in idea creativity between at least two groups ($F(3, 596) = 16.69$, $p = 2.01e-10$). Tukey's HSD test for multiple comparisons found that the mean value of idea quality was significantly different between the adaptive condition and the random condition, but no significant differences were found between the neutral and random conditions (see Fig. 4). The model can also adapt quickly; there is no significant difference between the condition where the transition function is pre-trained with random conversations and the simple adaptif model.

5 CONCLUSION AND FUTURE WORKS

In this work, we develop a scenario for a social robot to help with a creative task, such as an alternative use task. Our aim is to study the impact of social strategies on user mood and creativity. We also developed a POMDP model with a learnable transition function to adapt to each participant. The model was tested in interaction simulated users with promising results. In future work, we intend to replicate this experiment with human participants. The agent will be implemented using a Furhat robot [21], and the user's emotion will be detected using a fusion of state-of-the-art models on both speech and facial expression.

The goal of the evaluation will be to determine whether social strategies improve the acceptance of robot input, and thus creativity, and whether adapting to users' mood improves real users' creativity. Evaluation results can also be studied to determine whether different types of user can be computed and to study whether the link between mood and creativity is the same for all users. The results of the evaluation can also be used to develop a simulated user based on the data for this task.

REFERENCES

[1] Amin G. Alhashim, Megan Marshall, Tess Hartog, Rafal Jonczyk, Danielle Dickson, Janet van Hell, GÃ¼l E. Okudan-Kremer, and Zahed Siddique. 2020. WIP: Assessing creativity of alternative uses task responses: A detailed procedure. *ASEE Annual Conference and Exposition, Conference Proceedings* 2020-June.

[2] Safinah Ali, Nisha Devasia, Hae Won Park, and Cynthia Breazeal. 2021. Social Robots as Creativity Eliciting Agents. *Frontiers in Robotics and AI* 8 (2021).

[3] Patrícia Alves-Oliveira, Patrícia Arriaga, Matthew A Cronin, and Ana Paiva. 2020. Creativity encounters between children and robots. In *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 379–388.

[4] Teresa M Amabile, Sigal G Barsade, Jennifer S Mueller, and Barry M Staw. 2005. Affect and creativity at work. *Administrative science quarterly* 50, 3 (2005), 367–403.

[5] Roger E Beaty and Dan R Johnson. 2021. Automating creativity assessment with SemDis: An open platform for computing semantic distance. *Behavior research methods* 53, 2 (2021), 757–780.

[6] Alwin de Rooij, Philip J Corr, and Sara Jones. 2017. Creativity and emotion: Enhancing creative thinking by the manipulation of computational feedback to determine emotional intensity. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 148–157.

[7] Patrick Gebhard. 2005. ALMA: a layered model of affect. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*. 29–36.

[8] Milene Gonçalves, Carlos Cardoso, and Petra Badke-Schaub. 2016. Inspiration choices that matter: the selection of external stimuli during ideation. *Design Science* 2 (2016).

[9] Milene Gonçalves and Philip Cash. 2021. The life cycle of creative ideas: Towards a dual-process theory of ideation. *Design Studies* 72 (2021), 100988.

[10] Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. (2017). To appear.

[11] Yaxin Hu, Lingjie Feng, Bilge Mutlu, and Henny Admoni. 2021. Exploring the Role of Social Robot Behaviors in a Creative Activity. In *Designing Interactive Systems Conference 2021*. 1380–1389.

[12] Angel Hsing-Chi Hwang and Andrea Stevenson Won. 2021. IdeaBot: Investigating Social Facilitation in Human-Machine Team Creativity. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–16.

[13] Bahar Irfan, Anika Narayanan, and James Kennedy. 2020. Dynamic emotional language adaptation in multiparty interactions with agents. In *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*. 1–8.

[14] Alice M Isen, Mitzi Johnson, Elizabeth Mertz, and Gregory F Robinson. 1985. The influence of positive affect on the unusualness of word associations. *Journal of personality and social psychology* 48, 6 (1985), 1413.

[15] Peter H Kahn, Takayuki Kanda, Hiroshi Ishiguro, Brian T Gill, Solace Shen, Jolina H Ruckert, and Heather E Gary. 2016. Human creativity can be facilitated through interacting with a social robot. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 173–180.

[16] Joao Laranjeira, Goreti Marreiros, Carlos Freitas, Ricardo Santos, Joao Carneiro, and Carlos Ramos. 2011. A proposed model to include social and emotional context in a Group Idea Generation Support System. In *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*. IEEE, 609–612.

[17] Matthew V Law, JiHyun Jeong, Amritansh Kwatra, Malte F Jung, and Guy Hoffman. 2019. Negotiating the creative space in human-robot collaborative design. In *Proceedings of the 2019 on Designing Interactive Systems Conference*. 645–657.

[18] Tony Manstead. 2005. The social dimension of emotion. *the Psychologist* (2005).

[19] Gonçalo S Martins, Hend Al Tair, Luís Santos, and Jorge Dias. 2019. αPOMDP: POMDP-based user-adaptive decision-making for social robots. *Pattern Recognition Letters* 118 (2019), 94–103.

[20] William N Morris. 1989. Theories of mood. In *Mood*. Springer, 11–18.

[21] Samer Al Moubayed, Gabriel Skantze, and Jonas Beskow. 2013. The furhat back-projected humanoid head-lip reading, gaze and multi-party interaction. *International Journal of Humanoid Robotics* 10 (2013). Issue 1. https://doi.org/10.1142/S0219843613500059

[22] Mark A Runco and Garrett J Jaeger. 2012. The standard definition of creativity. *Creativity research journal* 24, 1 (2012), 92–96.

[23] Ben Shneiderman. 2007. Creativity support tools: Accelerating discovery and innovation. *Commun. ACM* 50, 12 (2007), 20–32.

[24] Robert J Sternberg and Todd I Lubart. 1999. The concept of creativity: Prospects and paradigms. *Handbook of creativity* 1, 3-15 (1999).

[25] Bas R Steunebrink, Mehdi Dastani, and John-Jules Ch Meyer. 2009. The OCC model revisited. In *Proc. of the 4th Workshop on Emotion and Computing*. Association for the Advancement of Artificial Intelligence Palo Alto, 62.

[26] Oshin Vartanian, Ingrid Smith, Timothy K Lam, Kristen King, Quan Lam, and Erin L Beatty. 2020. The relationship between methods of scoring the alternate uses task and the neural correlates of divergent thinking: Evidence from voxel-based morphometry. *NeuroImage* 223 (2020), 117325.

[27] Robert Weisberg. 1993. *Creativity: Beyond the myth of genius.* WH Freeman.

[28] Christina Wiethof, Navid Tavanapour, and Eva Bittner. 2021. Implementing an intelligent collaborative agent as teammate in collaborative writing: toward a synergy of humans and AI. In *Proceedings of the 54th Hawaii International*

Conference on System Sciences. 400.
[29] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. https://www.aclweb.org/anthology/2020.emnlp-demos.6
[30] Jianhua Yin and Jianyong Wang. 2014. A Dirichlet multinomial mixture model-based approach for short text clustering. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. https://doi.org/10.1145/2623330.2623715