TRIP NEGOTIATOR: A Travel Persona-aware Reinforced Dialogue Generation Model for Personalized Integrative Negotiation in Tourism

Priyanshu Priya[†], **Desai Vishesh Yasheshbhai**[†], **Ratnesh Kumar Joshi**[†], **Roshni Ramnani**[‡], **Anutosh Maitra**[‡], **Shubhashis Sengupta**[‡], **Asif Ekbal**[†]
[†]Department of Computer Science and Engineering, Indian Institute of Technology Patna, India
[‡]Accenture Labs, Banglore, India
[†]{priyanshu_2021cs26,desai_2211ai07,ratnesh_1921cs28,asif}@iitp.ac.in
[‡]{roshni.r.ramnani,anutosh.maitra,shubhashis.sengupta}@accenture.com

Abstract

A sophisticated negotiation dialogue system for tourism should engage in negotiations beyond mere price considerations, encompassing various other aspects and amenities inherent in the tourism package. To ensure such tailored interaction, it is imperative to understand the intricacies of traveler preferences, constraints, and expectations. Incorporating these personality facets allows for customizing negotiation strategies, resulting in a more personalized and integrative experience. With this aim, we take a pivotal step in advancing automated dialogue systems for personalized integrative negotiation tasks. We develop DEAL, a pioneering **D**ialogue datas**E**t for person**AL**ized integrative negotiation task in the tourism domain. Further, we propose TRIP NEGOTIATOR, a novel **T**ravel persona-aware **R**einforced d**I**alogue generation model for **P**ersonalized i**N**tegrative n**EGOTIAT**ion within the t**OuR**ism domain. TRIP NEGOTIATOR is built to discern the traveler's persona and intent, systematically adjusts negotiation strategies, and directs the negotiation toward a pertinent phase to ensure effective negotiation. Through reinforcement learning with Proximal Policy Optimization (PPO), we guide TRIP NEGOTIATOR to generate coherent and diverse responses consistent with the traveler's personality. Extensive qualitative and quantitative analyses demonstrate the effectiveness of TRIP NEGOTIATOR in generating personalized responses during negotiation[1].

1 Introduction

E-tourism's substantial growth over the past decade (Lazăr et al., 2019) has increased online trip bookings, necessitating the development of automated negotiation dialogue systems (NDS) for travelers and travel agencies. The tourism industry must comprehend travelers' *travel persona* to interpret their travel preferences and facilitate personalized interactions (Park et al., 2010). Besides, it needs to establish trust and rapport with travelers to foster collaborative and long-term relationships (Assaker and Hallak, 2013). Research has shown that modeling users' personality leads to successful negotiation (Yang et al., 2020), and an integrative approach to negotiation further elevates negotiation by promoting user inclusivity in the commercial arena (Guttman and Maes, 1998). Existing NDS (Dhingra et al., 2017; He et al., 2018; Zhou et al., 2019) fall short in modeling users' detailed persona encompassing their likes, dislikes, or priorities toward certain options, which hinders the generation of tailored strategies and utterances for diverse users. To address this gap, we attempt to consolidate personality and strategic modeling in NDS and propose a novel task of personalized integrative negotiation within the tourism domain. Personalized integrative negotiation aims to integrate travelers' travel persona and needs to ensure personally appealing and mutually beneficial negotiation consequences.

Travel persona defines travelers' preferences regarding a trip, spanning both fundamental (hotel types, meals, etc.) and optional (entertainment, outdoor activities, etc.) aspects and amenities. Beyond primary price considerations, travelers are interested in distinct fundamental aspects and optional amenities to varying extents. The negotiation on these factors depends on their interests. For instance, in Figure 1, a traveler interested in adventure negotiates for packages including activities like snowboarding and nature photography, while a mountain enthusiast emphasizes visits to scenic viewpoints and enjoying authentic local food. The agent's recommendations aligned with travelers' likes, dislikes, and priorities ensure a personalized experience and contribute to successful negotiation. Conversely, failing to capture this information leads to generic or unsuitable suggestions that can frustrate travelers and reduce the likelihood of a

[1]Dataset and code can be accessed at https://www.iitp.ac.in/~ai-nlp-ml/resources.html; https://github.com/priyanshu-profile/TRIP-NEGOTIATOR/.

Findings of the Association for Computational Linguistics: EACL 2024, pages 16566–16595

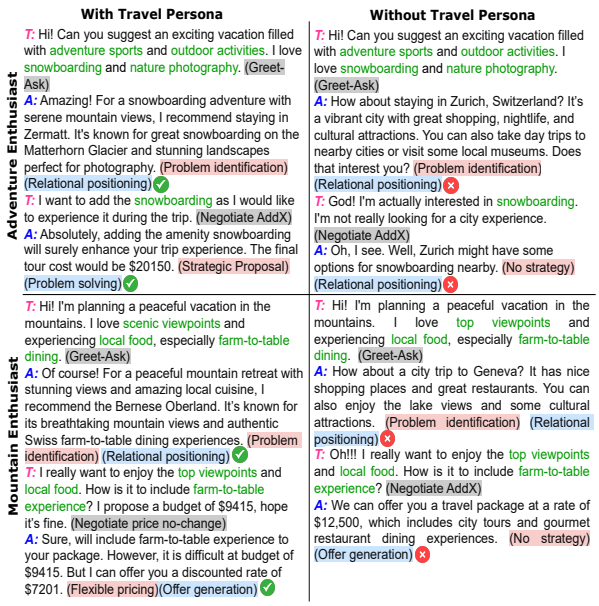

Figure 1: Example illustrating how personality and strategic modeling ensures personalized integrative negotiation and overlooking these facets results in negotiation failure. Green and grey text depict the traveler's persona and intent, while pink and blue text indicate the agent's negotiation strategy and phase, respectively. Here, *T* is Traveler and *A* is Travel agent.

successful booking.

The negotiation dialogues exhibit unique peculiarities like proposing prices, offering proposals, accepting or rejecting them, and/or formulating counter-offers (Thompson et al., 2010). Hence, during negotiation, the travel agent must employ an appropriate negotiation strategy as per the conversation context to resolve conflicts and foster a collaborative ambience. Further, the travel agent should discern the travelers' specific goals (intents) to recognize the key signals (proposals, acceptances, rejections, or counter-offers) and understand their requirements, eventually navigating the conversation toward an adequate negotiation phase. To exemplify, in Figure 1, the agent adapts negotiation strategies based on the context. Also, based on the traveler's intent to add an amenity or negotiate on price, the agent smoothly transitions the conversation to the relevant phase (problem-solving or offer-generation). The inability to do so leads to traveler dissatisfaction and negotiation breakdown.

Driven by these considerations, in this work, we propose TRIP NEGOTIATOR, a novel **T**ravel persona-aware **R**einforced d**I**alogue generation model for **P**ersonalized i**N**tegrative n**EGOTIAT**ion within the t**OuR**ism domain. To develop TRIP NEGOTIATOR, we first curate DEAL, a pioneering **D**ialogue datas**E**t for person**AL**ized integrative negotiation. DEAL is created through prompting Large Language Model (LLM), MPT-7B (Team et al., 2023) under few-shot settings followed by human interventions to ensure high-quality dialogues. The traveler's utterances in DEAL are then annotated with travel personas and intents, and the agent's utterances are annotated with negotiation strategies and phases. Using DEAL, TRIP NEGOTIATOR is built upon the DialoGPT (Zhang et al., 2020a) in a reinforcement learning (RL) framework. To train the end-to-end TRIP NEGOTIATOR, we use a novel reward function to ensure that the responses align with the traveler's persona and intent and adhere to the appropriate negotiation strategy and phase while maintaining fluency, contextual coherence, and diversity. Finally, we optimize a policy via RL using the Proximal Policy Optimization (PPO) algorithm (Schulman et al., 2017) based on the reward assigned to the generated responses.

In conclusion, the key contributions of this work are: **(i)** Introduce a novel task of personalized integrative negotiation within the tourism domain, which considers traveler's unique travel persona and needs to generate personalized integrative responses; **(ii)** Curate a novel dialogue dataset, DEAL through prompting for personalized integrative negotiation; **(iii)** Annotate DEAL with travel personas, negotiation strategies, intents, and negotiation phases; **(iv)** Propose an RL-based robust dialogue system, TRIP NEGOTIATOR to generate personalized integrative responses during negotiation focused on tourism domain by devising a pioneering reward function; **(v)** Conduct thorough automatic and human evaluations on the DEAL dataset to establish the efficacy of the TRIP NEGOTIATOR.

2 Related Work

Negotiation Dialogue Datasets. The development of negotiation dialogue datasets is crucial for advancing research in NDS. Existing datasets for distributive negotiation, such as CraigslistBargain (He et al., 2018), NegoCoach (Zhou et al., 2019), AntiScam (Li et al., 2020), and for integrative negotiation, like STAC (Asher et al., 2016), DealorNoDeal (Lewis et al., 2017), CaSiNo (Chawla et al., 2021), JobInterview (Yamaguchi et al., 2021), DinG (Boritchev and Amblard, 2022) have provided foundational insights into human negotiation behavior. These datasets focus on specific scenar-

ios, such as price bargaining, privacy protection, strategic games, item assignment, or job interviews, featuring simplified negotiation scenarios that may not fully capture the nuances of real-world negotiations. A recent dataset, IND (Ahmad et al., 2023), has introduced richer, more diverse integrative negotiation scenarios, including negotiation on price and other factors like item addition/removal from a deal bundle in an e-commerce setup. In contrast to these datasets, we introduce a novel dialogue dataset for the tourism domain, comprising integrative negotiation conversations on varied aspects and amenities of tourism packages, like price, destination, entertainment, transportation, and more.

Negotiation Dialogue Systems. Research in NDS has seen significant progress in recent years (Yamaguchi et al., 2021; Chawla et al., 2022; Fu et al., 2023; Deng et al., 2023b). Zhou et al. (2019); Deng et al. (2023a) proposed a distributive strategic approach aimed at recommending strategies to sellers to improve their deal outcomes. Zhao et al. (2019); Yamaguchi et al. (2021) introduced a collaborative strategy framework to mutually negotiate on the target issue. He et al. (2018) proposed dialogue acts tracking to model the behavior of buyers and sellers. Yang et al. (2020) improved the NDS by capturing the partners' personalities using a one-step dialogue-act look ahead during negotiation. Zhang et al. (2020b); Chawla et al. (2022) presented an opponent behavior modeling method to estimate opponent action during negotiation. Recently, Ahmad et al. (2023) introduced an integrative negotiation agent (INA) to negotiate item prices and tailor deals to customers' needs.

Existing research emphasized solely the objective outcomes of negotiation, such as points scored or the final agreed price. Also, the limited understanding of individual differences to accommodate varying preferences in these works restricts high-quality personalized conversations. Unlike these studies, the current work presents a novel personalized integrative negotiation system for tourism. This system utilizes the travelers' detailed travel persona to grasp their trip preferences for generating personally relevant responses. Our work further differentiates in that we propose novel integrative negotiation strategies to assist the agent in determining the next action based on the negotiation context and guiding the negotiation toward the desired phase whilst aligned with the travelers' intent. Following (Shi et al., 2021; Ahmad et al., 2023), we train our system with standard PPO loss using six novel rewards customized for the proposed task. To our comprehension, our work pioneers the development of such a personalized integrative NDS for the tourism domain.

3 Dataset

To develop the personalized trip negotiator, we crafted DEAL, an innovative **D**ialogue datas**E**t for person**AL**ized integrative negotiation task. We focus on the dialogues considering the traveler's persona information and requirements to facilitate the best travel packages and deals. The ultimate goal is to offer a seamless and collaborative planning and booking experience during negotiation, thereby contributing to an enriched and positively transformed landscape within the realm of tourism.

3.1 Dataset Creation

DEAL comprises interactions between the travel agent and the traveler negotiating on price, destination, various aspects, and amenities of the tourism package. To alleviate the reliance on expensive human resources, the dataset is developed by leveraging the extensive knowledge embedded in the LLM, MPT-7B (Team et al., 2023). Specifically, the dataset is created by prompting the MPT-7B model, followed by human intervention to ensure quality control. The entire dataset creation process encompasses *four* key stages: (a) Background Data Acquisition, (b) Negotiation-specific Intent Formulation, (c) Dialogue Flow Generation, (d) Traveler-Agent Dialogue Curation.

(a) Background Data Acquisition. This stage involves gathering relevant details about diverse facets of the tourism domain, such as tourist destinations, attractions, accommodations, transportation options, and other related factors. This information helps formulate 20 distinct travel packages, each accompanied by various aspects, amenities, and services to cater to diverse preferences and needs. Overall, the background database used to create dialogues comprises essential details like package names and their description, list of aspects, amenities, and services with their descriptions. The list of packages is given in §A.1 of the appendix.

(b) Negotiation-specific Intent Formulation. To build robust integrative negotiation dialogue systems, it is crucial to formulate intents encompassing a broad spectrum of negotiation scenarios, including price and feature-based negotiation. To comprehensively address these diverse scenarios,

we formulate the 17 intents, namely *Greet, Ask, Inform, Ask price, Tell price, Ask clarification-Y, Provide clarification-Y, Negotiate price increase, Negotiate price decrease, Negotiate price no-change, Negotiate add-X, Negotiate remove-X, Provide consent, Consent response, Accept, Reject,* and *Acknowledge acceptance*. The definitions and examples of intents are detailed in Table 6 and Table 7, respectively, in the appendix.

(c) Dialogue Flow Generation. The dialogue flow is essentially a sequence of intents occurring during negotiation. We devise a *Dialogue Flow Generator* (DFG) module to automatically generate the dialogue flows. The DFG module assumes that dialogue flow can exhibit randomness; however, to ensure consistency in generated dialogues, we impose a constraint of initiating a dialogue by the traveler with a 'Greet' intent, which may be followed by a clarification request or one of the intents designed to negotiate on certain attributes (price, aspects/amenities). The travel agent can then respond by the 'Inform' intent or one of the intents demonstrating negotiation on distinct attributes.

We also maintain negotiation information on the ongoing deal with the traveler for each travel package. This information comprises the agent's minimum price and current price, the traveler's current price, and the tolerance threshold (ψ). To uphold the win-win outcomes of the negotiation, we confine price-based negotiations to predefined number of dialogue turns, after which the intents 'Negotiate add-X' or 'Negotiate remove-X' come into play. Further, to propose the price for the subsequent utterance, we assume a decline in the price difference (increment for the traveler and decrement for the agent) across dialogue turns. This approach mirrors the method of Faratin et al. (1998), where a comparable function models price negotiations between customers and sellers. The proposed prices by the traveler (Pr_t) and agent (Pr_a) at the current turn, i, are computed as follows:

$$Pr_{a_i} = Pr_{t_{i-1}} + (Pr_{a_{i-1}} - Pr_{t_{i-1}})e^{-kt} \quad (1)$$

$$Pr_{t_i} = Pr_{a_{i-1}} + (Pr_{a_{i-1}} - Pr_{t_{i-1}})e^{-kt} \quad (2)$$

Here, k is a constant governing the rate of price change from one utterance to the next. A higher value of k results in a higher concession rate, whereas a lower value indicates a reduced concession rate from the agent. In our specific scenario, we assume a higher value of k for the agent and a lower k for the traveler, reflecting the traveler's tendency to adhere closely to their budget constraints. Also the value of k for agent depends upon the initial price gap between traveler and agent. For low-budget, where price gap is high, we choose higher value of k and for high-budget, where price gap is low, we choose lower value of k (c.f. Table 4 in appendix). During negotiation, the agent will opt for the 'Accept' intent if the offered price from the traveler $\leq (Pr_{a_i} - \psi * Pr_{a_i})$. Conversely, the traveler will select the 'Reject' intent when the negotiation deadline has been surpassed and the agent is unwilling to reduce the package price further. The dialogue concludes upon acknowledging with either the 'Accept' or 'Reject' intent.

(d) Traveler-Agent Dialogue Curation. The traveler-agent dialogue curation involves prompt designing and dialogue generation.

Prompt Designing. To begin the few-shot dialogue generation using the MPT-7B model, we design few-shot prompts (Brown et al., 2020) for each intent. Each few-shot prompt consists of three distinct components: a task description, a concise summary of the pertinent information from the dialogue, and an utterance aligned with the specified intent. The task description elucidates the intent's scenario and objective; for instance, the task description for the 'Acknowledge acceptance' intent might entail the traveler agreeing to opt for a package, causing the agent to express gratitude towards the traveler and proceed with the transaction. The summary of the pertinent information is devised considering the sequential flow of intents from the preceding dialogue utterances. The utterance aligned with the specified intent is manually created using the task description and the information summary of the few-shot prompt. A sample prompt is provided in §A.2 of the appendix.

Dialogue Generation. Subsequently, the DFG module generates an ordered list of intents accompanied by relevant details for each intent. For example, the proposed price adjustment is articulated for the intent 'Negotiate price decrease', while for 'Negotiate remove-X', the item to be removed is specified. We leverage this generated list of intents to augment the prompt corresponding to each intent. This augmented prompt is then used to prompt the MPT-7B model to generate the utterance. Figure 3 in appendix depicts entire dataset creation process.

3.2 Data Cleaning and Quality Control

After obtaining the entire dialogue dataset, each dialogue is manually evaluated to maintain the dataset's adherence to quality standards. For man-

ual verification, we recruit three human experts, two holding doctoral degrees in Linguistics and one holding a post-graduate degree in Computer Science. All experts possess two years of expertise in post-editing. We instruct them to review and post-edit the automatically generated dialogues to ensure grounding in the provided background data, intent, action, and negotiation dialogue flow. Further, we ask them to rate each utterance of the dialogues for fluency[2] and naturalness[3]. We observe an agreement ratio (McHugh, 2012) of 84.9% for fluency and 87.2% for naturalness among these experts. The dialogues containing the agent's utterances that reflect the agent's feelings or experiences, pretending to be human, are dropped to ensure authenticity. Besides, the dialogues with utterances rated 1 for fluency and 0 for naturalness are dropped from the dataset. The final dataset statistics are given in Table 1.

3.3 Dataset Quality Assessment

Automatic assessment. We automatically assess the quality of the proposed DEAL dataset following (Wang et al., 2021). In particular, we employ BLEU-1 (Papineni et al., 2002) and METEOR (Banerjee and Lavie, 2005) metrics to automatically assess the variability of the dialogues in the DEAL dataset. The low BLEU-1 and METEOR scores of 0.09 and 0.04, respectively, signify a high degree of variability between dialogue utterances in the dataset.

Manual assessment. After obtaining the entire conversational dataset, each dialogue is manually assessed in terms of *Naturalness* (Nat.), *Engagingness* (Eng.), *Fairness* (Fair.), and *Faithfulness* (Faith.) on a scale of 1-3 (1-low, 2-moderate, 3-high) by the same group of human experts involved in the data cleaning. *Naturalness* measures the degree to which the generated dialogues mimic the human conversation, *Engagingness* measures how well the generated dialogues maintain the user's interest throughout the conversation, *Fairness* assesses the equitability of the final negotiation outcome, gauging how well it leads to a win-win scenario, and *Faithfulness* assesses whether the generated dialogues align with the travelers' travel personas, i..e., travelers' preferences and priorities concerning various aspects, amenities, and services of the travel package. The dialogues obtain average ratings of 2.54, 2.78, 2.34, and 2.69 for Nat., Eng., Fair., and Faith., respectively. Moreover, an agreement ratio of 86.9%, 84.8%, 87.2%, and 79.8% for Nat., Eng., Fair., and Faith., respectively, is observed among these experts. Due to space constraints, details on the manual assessment of our LLM-generated dialogues with those of human-generated ones are given in §A.3 of the appendix.

Metrics	Train	Validation	Test
# of Dialogues	904	129	258
# of Utterances	13,207	1,866	3,859
Avg. utterances per dialogue	14.61	14.47	14.96
Avg # of words in traveler utterance	18.58	18.75	17.96
Avg # of words in travel agent utterance	28.36	27.44	26.60

Table 1: Dataset statistics of DEAL.

3.4 Dataset Annotation

The traveler's utterances in the DEAL dataset are annotated with 'travel persona' using *Persona Extraction Module* (PEM), while the agent's utterances are annotated with negotiation strategies and phases employing a semi-supervised approach.

Personalized trip negotiation dialogues require creating detailed, domain-dependent travel personas that reflect the traveler's destination choice, holiday priorities, and other travel-related decisions like accommodation, travel groups, and leisure activities. Hence, we leverage travel personas that focus on travelers' preferences and priorities, ensuring that their trip preferences are adequately met.

Further, at each step of negotiation, the travel agent often needs to employ an appropriate negotiation strategy to choose among different possible actions to achieve successful negotiations and foster collaborative outcomes. Hence, we devise a set of negotiation strategies for the travel agent using the negotiation theories and a preliminary assessment of 60 randomly selected dialogue samples. The three human experts[4] independently labeled sampled dialogues, analyzed discrepancies and modified the strategies accordingly. The final inter-expert agreement (McHugh, 2012) surpassed 0.83 for all categories. Specifically, *eight* different negotiation strategies, namely *problem identification (pi)*, *strategic proposal (sp)*, *firm pricing (fp)*, *definitive decision making (dd)*, *collaborative proposal (cp)*, *flexible pricing (flp)*, *co-operative decision making (cd)* and *no strategy (ns)* are identified.

[2] 1-5 (5: Flawless, 4: Good, 3: Non-native, 2: Disfluent, 1: Incomprehensible)

[3] 0-2 (2: high, 1: moderate, 0: low)

[4] Two possess Ph.Ds in Linguistics, and one has an M.Sc. in Computer Science. They have two years of experience in related tasks and are paid according to institute norms.

Also, negotiation typically unfolds through distinct phases or stages. Negotiation phases provide a narrative explanation of the negotiation process, i.e., they identify sequences of events that constitute the entire negotiation. Thus, to apprehend the negotiation progress from inception to conclusion, we devise the *four* negotiation phases, namely *relational positioning*, *problem solving*, *offer generation*, and *decision making*. Due to space constraints, definition and example utterances of negotiation strategies and phases are given in Table 8 and Table 9, respectively, and the annotation procedures are provided in §A.5 in the appendix.

4 Proposed TRIP NEGOTIATOR

To facilitate effective negotiation between the dialogue agent and travelers having unique personas and intentions, we introduce TRIP NEGOTIATOR, a novel **T**ravel persona-aware **R**einforced d**I**alogue generation model for **P**ersonalized i**N**tegrative n**EGOTIAT**ion within t**OU**r**R**ism domain. The architecture of TRIP NEGOTIATOR is depicted in Figure 2. It is built in three stages, *viz.* Supervised Fine-tuning, Reward Modeling, and RL Fine-tuning.

Supervised Fine-tuning Stage. Let, $\mathcal{D} = \{t_1, a_1, t_2, a_2, ..., t_T, a_T\}$ be a multi-turn dialogue between the traveler (t) and travel agent (a), where t_i and a_i denote the traveler's and the agent's i^{th} utterance, respectively, and T denote the total number of utterances in $\mathcal{D}$. Let the traveler's current utterance t_i with the corresponding travel persona $\mathcal{P}_i$ be represented as $t_{\mathcal{P}_i} = [t_i \oplus \mathcal{P}_i]$. Further, let π_a and $\pi_{t_\mathcal{P}}$ denote the probability distributions of the agent's utterance and the traveler's utterance with corresponding travel persona, respectively. Then, for a given context, a persona-aware supervised fine-tuned model (PSFT) is obtained by maximizing the likelihood over Equation 3.

$$\pi_{\text{PSFT}}(\mathcal{D}) = \prod_{i=1}^{T} \pi_{t_\mathcal{P}}(t_i | t_{\mathcal{P}_{<i}}, a_{<i}) \pi_a(a_i | t_{\mathcal{P}_{\leq i}}, a_{<i}) \quad (3)$$

Reward Modeling Stage. The reward model aims to generate persona-aligned responses that adhere to pertinent negotiation strategies and drive the negotiation toward the desired phase, considering the travelers' intents while maintaining fluency, contextual coherence, and diversity in the generated responses. To achieve this, a reward model comprising *Trait-focused Rewards* (TFR), *viz.* Persona Alignment, Negotiation Strategy Consistency, and Intent-Phase Agreement rewards, and *Language-focused Rewards* (LFR), *viz.* Response Fluency, Dialogue Coherence, and Diversity rewards is devised. TFR guides the agent toward the responses exhibiting pertinent negotiation strategy and phase while taking into account the travelers' persona and intent, and LFR ensures the syntactic and semantic richness of the generated responses.

Persona Alignment Reward (PAR). Aligning the negotiation agent's response to match the traveler's persona encourages personalized interactions. We approach the problem of response alignment with traveler's persona, $\mathcal{P} = \{p_1, p_2, ..., p_l\}$ as a Natural Language Inference (NLI) task defined as $f_{\text{NLI}}(\hat{r}_i, p_k) \rightarrow \{E, N, C\}$ (*E: entailment, N: neutral*, and *C: contradiction*). Entailed responses align with the persona, whereas contradictory responses are inconsistent and thus warrant penalization.

To assess the persona alignment of the generated responses, a BERT-based (Devlin et al., 2018) Persona-NLI model (Persona-NLIM) is first fine-tuned on Dialogue NLI dataset (Welleck et al., 2018)[5]. This fine-tuned model is then run on each $(\hat{r}_i, p_k)$ pair to obtain raw outputs as $\mathcal{O}_{i,k} = \{o^E_{i,k}, o^N_{i,k}, o^C_{i,k}\}$, which is used to attain the respective class probability as $\rho_{i,k} = softmax(\mathcal{O}_{i,k}) = \{\rho^E_{i,k}, \rho^N_{i,k}, \rho^C_{i,k}\}$. These probability scores are finally used to compute PAR as $\mathcal{R}_{pa} = \sum_{k=1}^{l}(1 - \rho^C_{i,k})$. It can be inferred that $\mathcal{R}_{pa}$ will be more when the generated response does not contradict the persona.

Negotiation Strategy Consistency Reward (NSCR). During negotiation, travel agents must employ a well-thought-out strategy to manage conflicts and ensure win-win outcomes. Thus, we penalize the responses deviating from the ground-truth negotiation strategy. To design NSCR, we track class probabilities score from the negotiation strategy classification model (NSCM) and compute NSCR as: $\mathcal{R}_{nsc} = NSCM(t_i) - \delta \times NSCM(\hat{r}_i)$.

NSCM is built by fine-tuning RoBERTa-large (Radford et al., 2019) on the DEAL dataset[6]. Here, $\delta \geq 1$ serves as the penalization factor.

Intent-Phase Agreement Reward (IPAR). The travel agent should steer the negotiation toward the decision phase, empowering travelers to make informed choices for a more satisfying and mutually beneficial outcome. With the evolving traveler's utterance and its intent, travel agent should negotiate differently and advance the conversation to the optimal phase. However, it is challenging to model the association between travelers's dynami-

[5] We attain an accuracy of 86.76% on the test set.

[6] We obtain an accuracy and macro-F1 of 91.75% and 81.34%, respectively.

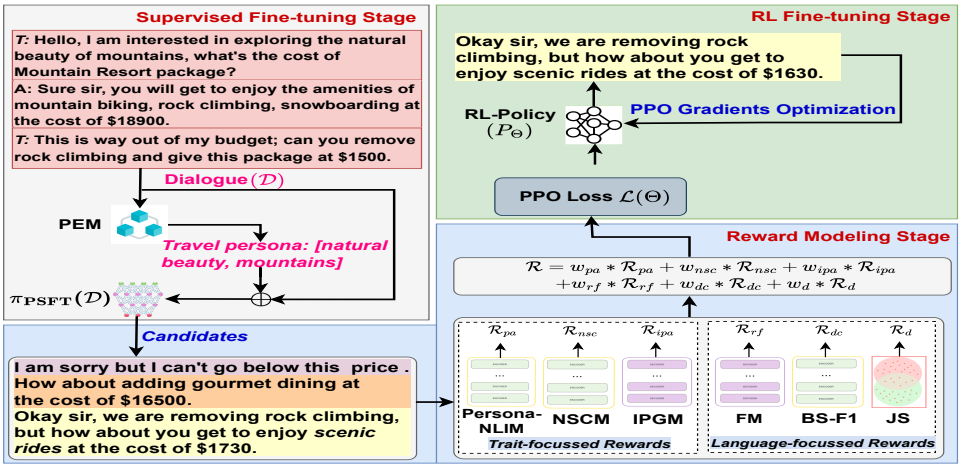

Figure 2: Overall architecture of personalized integrative negotiation dialogue system - TRIP NEGOTIATOR.

cally changing intents and the travel agent's ability to guide the conversation toward the intended phase. Hence, to enable the agent to employ intent-aware correct negotiation phase at i^{th} turn, we formulate the IPAR as $\mathcal{R}_{ipa} = \sigma(J(P(a_i, ph_i|t_i, in_i)) - \eta \times J(P(\hat{r}_i, \hat{ph}_i|t_i, in_i)))$, where, $P(\cdot)$ signifies the intent-phase generation model (IPGM) developed by fine-tuning DialoGPT-medium (Zhang et al., 2020a) on DEAL[7]. It takes traveler's utterance and its intent as input and generates travel agent's utterance and corresponding phase as output. in_i and ph_i represent the traveler's intent and agent's phase, respectively. $\hat{ph}_i$ denotes the generated response's phase predicted by the Negotiation Phase Classification Model (NPCM) obtained by fine-tuning RoBERTa-large[8]. $J(\cdot)$ denotes negative log-likelihood loss function and $\sigma(\cdot)$ denotes the $tanh()$[9] activation function, which scales down the values to $[0, 1]$.

Response Fluency Reward (RFR). The generated response should be fluent (grammatically correct). Hence, we design RFR as the inverse of the perplexity of generated response $\hat{r}_i$ as $\mathcal{R}_{rf} = \text{FM}(\hat{r}_i)^{1/m}$. We use PLM, DialoGPT-medium as fluency model (FM), and m is the word count in $\hat{r}_i$.

Dialogue Coherence Reward (DCR). The generated response should be coherent with the dialogue context to ensure a smooth flow of conversation. Hence, we devise the DCR, which penalizes the generated response ($\hat{r}_i$) deviating from the context (c_i) and the traveler's utterance (t_i) as $\mathcal{R}_{dc} = \text{MIN}([h(c_i, \hat{r}_i) + h(t_i, \hat{r}_i)], 1)/2$, where, $h(\cdot)$ denotes the similarity measure computed using BERTScore-F1 (BS-F1) (Zhang et al., 2019). To mitigate the occurrence of similar responses, a threshold value of 1 is employed.

Diversity Reward (DR). During negotiation, the generated response may become generic and repetitive (*e.g., I can offer the best package at best price*), which affects the overall conversation quality (Li et al., 2015; See et al., 2019). To ensure diverse and engaging responses, we formulate DR utilizing the Jaccard similarity (JS) between the responses, $\hat{r}_i$ and $\hat{r}_{i-1}$ at i^{th} and $(i-1)^{th}$ turns, respectively as $\mathcal{R}_d = 1 - (\hat{r}_{i-1} \cap \hat{r}_i)/(\hat{r}_{i-1} \cup \hat{r}_i)$.

Reward Function. To train the entire system, we formulate a reward function $\mathcal{R}$ as the weighted sum of all the rewards, i.e., $\mathcal{R} = w_{pa} * \mathcal{R}_{pa} + w_{nsc} * \mathcal{R}_{nsc} + w_{ipa} * \mathcal{R}_{ipa} + w_{rf} * \mathcal{R}_{rf} + w_{dc} * \mathcal{R}_{dc} + w_d * \mathcal{R}_d$. All reward values are between 0 and 1 and $\sum_{r \in W} w_r = 1$, where $W = \{w_{pa}, w_{nsc}, w_{ipa}, w_{rf}, w_{dc}, w_d\}$ is a set of weights (hyperparameters) that are tuned to optimize and maximize the reward.

RL Fine-tuning Stage. The final reward $\mathcal{R}$ is employed within an RL policy loss PPO, which offers feedback to the agent, motivating it to produce high-quality responses corresponding to the desired outcomes. During the RL fine-tuning phase, the RL policy P_Θ is initialized π_{PSFT}. Due to space constraints, training details and RL policy optimization with PPO are furnished in §B of the appendix.

5 Experimental Details

Baselines. We evaluate TRIP NEGOTIATOR against eight baselines, *viz.* GPT-2-small (Radford et al., 2019), DialoGPT-small (Zhang et al., 2020a), ARDM (Wu et al., 2021), GPT-Critic (Jang et al.,

[7] We attain PPL, BLEU, METEOR, and BS-F1 as 3.71, 0.29, 0.63, and 0.82, respectively.

[8] We attain an accuracy and macro-F1 of 93.24% and 84.02%, respectively.

[9] The $tanh()$ maps values between -1 and 1, but in this case, the loss cannot be negative, so the resulting values will always be between 0 and 1. Further, $sigmoid()$ is not chosen as it would map values between 0.5 and 1.

	Models	PPL	BLEU	METEOR	BS-F1	PA	NSC	NPA	R-LEN
Baselines	**GPT-2** (Radford et al., 2019)	5.90	0.18	0.58	0.62	64.52%	71.24%	68.72%	18.61
	DialoGPT (Zhang et al., 2020a)	5.27	0.20	0.59	0.64	66.83%	73.47%	69.23%	19.82
	ARDM (Wu et al., 2021)	4.82	0.23	0.65	0.69	71.25%	75.34%	72.46%	24.92
	PersRFI (Shi et al., 2021)	4.26	0.25	0.65	0.70	71.69%	75.05%	71.46%	25.03
	GPT-Critic (Jang et al., 2022)	4.12	0.24	0.67	0.70	71.01%	75.60%	71.69%	23.89
	INA (Ahmad et al., 2023)	3.98	0.27	0.69	0.71	70.28%	74.10%	70.15%	25.19
	ProCoT + Llama-2-7b-chat (Deng et al., 2023a)	36.87	0.11	0.33	0.29	36.14%	39.90%	32.86%	10.75
	ProCoT + ChatGPT (Deng et al., 2023a)	29.30	0.19	0.50	0.43	42.76%	46.30%	39.42%	19.46
	PSFT	3.67	0.28	0.75	0.73	74.58%	78.60%	74.36%	28.16
Proposed System	TRIP NEGOTIATOR	**2.13**	**0.35**	**0.87**	**0.81**	**82.31%**	**84.62%**	**81.34%**	**31.84**
Ablations	TRIP NEGOTIATOR - $\mathcal{R}_{pa}$	2.51	0.32	0.83	0.78	79.62%	82.61%	79.82%	29.63
	TRIP NEGOTIATOR - $\mathcal{R}_{nsc}$	2.46	0.33	0.85	0.80	80.16%	81.37%	80.36%	30.43
	TRIP NEGOTIATOR - $\mathcal{R}_{ipa}$	2.48	0.33	0.85	0.79	81.04%	82.99%	78.01%	30.19
	TRIP NEGOTIATOR - $\mathcal{R}_{rf}$	2.38	0.34	0.86	0.80	82.03%	84.16%	80.73%	30.72
	TRIP NEGOTIATOR - $\mathcal{R}_{dc}$	2.41	0.33	0.85	0.79	81.27%	83.28%	80.54%	30.76
	TRIP NEGOTIATOR - $\mathcal{R}_{d}$	2.37	0.34	0.85	0.80	81.64%	83.54%	80.59%	30.83
	TRIP NEGOTIATOR - ($\mathcal{R}_{pa}$ + $\mathcal{R}_{nsc}$ + $\mathcal{R}_{ipa}$)	2.58	0.31	0.81	0.77	77.24%	79.82%	78.52%	28.07
	TRIP NEGOTIATOR - ($\mathcal{R}_{rf}$ + $\mathcal{R}_{dc}$ + $\mathcal{R}_{d}$)	2.46	0.34	0.85	0.78	81.76%	84.03%	80.65%	30.72
	TRIP NEGOTIATOR - $\mathcal{R}$	3.06	0.29	0.76	0.73	73.10%	77.92%	74.81%	29.17
ΔTRIP NEGOTIATOR $- BEST_{Baseline}$		-41.96% ↓	25.00% ↑	16.00% ↑	10.95% ↑	10.36% ↑	7.65% ↑	9.38% ↑	20.17% ↑

Table 2: Automatic evaluation results of the TRIP NEGOTIATOR system on the proposed DEAL dataset. All results are statistically significant based on Welch's t-test (Welch, 1947) conducted at a 5% significance level.

	Models	PA	NSC	NPA	WT	DR	PO	F	CC	E
Baselines	**GPT-2** (Radford et al., 2019)	2.13	2.32	2.18	53.75%	58.44%	60.23%	2.05	2.16	1.87
	DialoGPT (Zhang et al., 2020a)	2.48	2.86	2.53	56.25%	61.39%	62.01%	2.34	2.43	2.08
	ARDM (Wu et al., 2021)	2.76	2.63	2.56	67.50%	65.99%	67.48%	2.67	2.7	2.18
	PersRFI (Shi et al., 2021)	2.93	2.78	2.72	70.25%	67.21%	68.50%	2.82	2.85	2.24
	GPT-Critic (Jang et al., 2022)	2.86	2.71	2.69	70.00%	66.80%	68.33%	2.74	2.79	2.22
	INA (Ahmad et al., 2023)	2.04	2.88	2.79	71.10%	68.00%	69.05%	2.92	2.91	2.26
	ProCoT + ChatGPT (Deng et al., 2023a)	1.80	1.56	1.49	30.12%	36.29%	35.01%	3.50	3.60	1.20
	PSFT	3.61	3.78	3.91	75.00%	71.32%	72.48%	3.91	3.55	3.24
Proposed System	TRIP NEGOTIATOR	**4.32**	**4.21**	**4.53**	**88.75%**	**82.76%**	**87.24%**	**4.73**	**4.71**	**4.19**
Ablations	TRIP NEGOTIATOR - ($\mathcal{R}_{pa}$ + $\mathcal{R}_{nsc}$ + $\mathcal{R}_{ipa}$)	3.86	3.81	4.12	77.50%	80.14%	84.17%	4.2	4.14	3.71
	TRIP NEGOTIATOR - ($\mathcal{R}_{rf}$ + $\mathcal{R}_{dc}$ + $\mathcal{R}_{d}$)	3.95	3.89	4.38	81.25	81.52	85.49	4.33	4.2	3.86
	TRIP NEGOTIATOR - $\mathcal{R}$	3.64	3.55	3.97	76.25	78.23	84.46	3.99	3.92	3.61

Table 3: Human evaluation results of the TRIP NEGOTIATOR on the proposed DEAL dataset. All results are statistically significant based on Welch's t-test (Welch, 1947) conducted at a 5% significance level.

2022), PersRFI (Shi et al., 2021), INA(Ahmad et al., 2023), ProCoT (Deng et al., 2023a), and PSFT: DialoGPT fine-tuned in a supervised fashion with traveler's persona. Baselines (§C.1) and implementation details (§C.2) are given in appendix.

Evaluation Metrics. We conduct automatic and human evaluations to assess the performance of the TRIP NEGOTIATOR. We evaluate classifiers using Weighted Accuracy (W-Acc.) and Macro-F1, and intent-phase generation model using Perplexity (PPL) (Brown et al., 1992), BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), and BERTScore-F1 (BS-F1) (Zhang et al., 2019).

TRIP NEGOTIATOR is evaluated using *trait-focused metrics* (TFM) to measure the travel agent's proximity to the goal from the perspective of persona modeling and negotiation, and *language-focused metrics* (LFM) to emphasize the language realization and generation ability for negotiation dialogue. For automatic evaluation, Persona Alignment (PA), Negotiation Strategy Consistency (NSC), and Negotiation Phase Agreement (NPA)[10] are used as TFM, and PPL, BLEU, METEOR, BS-F1, and response length (R-LEN) are used as LFM. For human evaluation, besides PA, NSC and NPA, Win Times (WT), Deal Rate (DR), and Pareto Optimality (PO) (Lewis et al., 2017) are utilized as TFM, and Fluency (F), Contextual Coherence (CC), and Engagingness (E) are used as LFM. PA, NSC, NPA, F, CC, and E are rated on a scale of 1-5 (the higher, the better). The details of evaluation metrics and human evaluation process are given in §C.5 of the appendix.

6 Results and Analysis

Automatic Evaluation Results. Table 2 presents the results of TRIP NEGOTIATOR and baselines. The superior performance of PSFT compared to other baselines, *viz.* GPT-2, DialoGPT, ARDM, PersRFI, GPT-Critic, INA, and LLMs' prompting-

[10] PA, NSC, NPA are computed using Persona-NLIM, NSCM, and NPCM accuracy scores.

based ProCoT underscores the significance of the travel persona. The integration of these attributes inherently guides PSFT towards generating more personalized and interactive responses. Further, the evaluation results show that TRIP NEGOTIATOR obtains the best scores compared to all baselines across all the metrics. Evidently, there is a significant decline in PPL. The proposed system obtains PPL of 2.13, which is -41.96% compared to the best baseline (PSFT). TRIP NEGOTIATOR also generates longer responses, which can be inferred from the highest score of R-LEN. This could be attributed to the TFR, RFR, and DCR that drive the model to build a connection with travelers by generating personalized integrative responses tailored to travelers' preferences and requirements while being contextually adequate and fluent. This results in the generation of interactive and engaging responses. The proposed model also shows an improvement in TFM. TRIP NEGOTIATOR reports improved PA, NSC, and NPA scores of 82.31%, 84.62%, and 81.34%, respectively, which are 10.36%, 7.65%, and 9.38% better than PSFT. This could be due to the TFR that forces the model to generate persona-aligned responses that adhere to pertinent negotiation strategies and drive the negotiation toward the desired phase based on the traveler's intent.

Further, the proposed system yields 0.35, 0.87, and 0.81 scores of BLEU, METEOR, and BS-F1, respectively, with an increase of +0.07, +0.12, and +0.08 points compared to PSFT. These scores signify high lexical and semantic similarity of generated responses with the gold responses. This strengthens the design of rewards in generating customized responses aligned with gold responses. It can also be deduced that the performance of TRIP NEGOTIATOR decreases when any of the rewards are removed. Notably, the PPL score increases in the order TRIP NEGOTIATOR - ($\mathcal{R}_{pa} < \mathcal{R}_{ipa} < \mathcal{R}_{nsc} < \mathcal{R}_{dc} < \mathcal{R}_{rf} < \mathcal{R}_{d}$). The performance further drops when either all the TFR, or LFR, or both are ablated. These results signify the need for both TFR and LFR to generate fluent, coherent, and diverse responses grounded in appropriate negotiation strategy and phase and aligned with the traveler's persona and intents. Due to space constraints, results of classifiers and intent-phase model are given in §C.3 and §C.4 of the appendix.

Human Evaluation Results. Table 3 presents human evaluation results for TRIP NEGOTIATOR. We compare our proposed model against GPT-2, DialoGPT, ARDM, PersRFI, GPT-Critic, INA, ProCoT, PSFT, TRIP NEGOTIATOR$-(\mathcal{R}_{pa}+\mathcal{R}_{nsc}+\mathcal{R}_{ipa})$, TRIP NEGOTIATOR$-(\mathcal{R}_{rf}+\mathcal{R}_{dc}+\mathcal{R}_{d})$ and TRIP NEGOTIATOR $-\mathcal{R}$ only as manual evaluation is expensive. It is evident that TRIP NEGOTIATOR achieves better scores of 4.32, 4.21, 4.53, 88.75%, 82.76%, 87.24%, 4.73, 4.71, and 4.19 for PA, NSC, NPA, WT, DR, PO, F, CC, and E, respectively, with an improvement of +0.71, +0.43, +0.62, +13.75, +11.44, +14.76, +0.82, +1.16, and +0.95, points for these metrics, compared to PSFT. The highest PA, NSC, NPA, WT, DR, and PO scores emphasize that incorporating PAR, NSCR, and IPAR helps TRIP NEGOTIATOR to offer personalized experience during negotiation by generating engaging and interactive responses.

The highest scores of F, CC, and E suggest that FR, DCR, and DR have played a crucial role in generating consistent, fluent, and non-repetitive responses. Also, TRIP NEGOTIATOR$-(\mathcal{R}_{pa}+\mathcal{R}_{nsc}+\mathcal{R}_{ipa})$ shows a marginal improvement on PSFT and TRIP NEGOTIATOR$-\mathcal{R}$. This further implies that LFR are insufficient to ensure persona-aware responses in negotiation conversations. Notably, a minimal performance difference can be seen between TRIP NEGOTIATOR and TRIP NEGOTIATOR $-(\mathcal{R}_{pa}+\mathcal{R}_{nsc}+\mathcal{R}_{ipa})$. This signifies that TFR are crucial for generating persona-aware responses with the correct negotiation strategy and phase. It can also be noted that TRIP NEGOTIATOR achieves better scores than TRIP NEGOTIATOR$-(\mathcal{R}_{rf}+\mathcal{R}_{dc}+\mathcal{R}_{d})$ for all metrics, indicating the importance of LFR.

7 Conclusion

This work introduces a new task of personalized integrative response generation considering the travelers' *travel persona* and requirements in negotiation conversations. For this task, a novel dialogue dataset DEAL is created by prompting MPT-7B with manual interventions and is annotated with travel persona, intents, negotiation strategies, and phases. Then, a novel personalized integrative negotiation dialogue system, TRIP NEGOTIATOR, is developed in an RL framework by designing unique rewards to facilitate the generation of travel persona-aware responses with appropriate negotiation strategy and negotiation phase aligned with the pertinent intent. The findings emphasize the necessity of each reward to facilitate a personalized integrative experience during negotiation.

Limitations

The present work has limitations regarding the proposed DEAL dataset creation approach and modeling of the proposed TRIP NEGOTIATOR framework. Regarding data curation, using the MPT-7B model for dialogue generation requires substantial GPU memory, specifically 40 GB. Another constraint arises from MPT-7B's context window, which is limited to 2,048 tokens. This restriction impacts our prompting mechanism, as within this window, we must accommodate background data alongside dialogue history and few-shot examples. Consequently, we are constrained to a maximum of 4 shots during prompting, leading to some hallucinations in the generated data, thus requiring manual correction.

Regarding the dialogue system, TRIP NEGOTIATOR also, one limitation is the requirement of substantial GPU memory capacity, specifically 40 GB for training the end-to-end RL model. Another challenge stems from optimizing reward weights, potentially extending both training and validation times. To mitigate this, heuristic approaches are employed to select specific combinations of reward weights. Additionally, in cases involving continuous, brief, or direct responses (e.g., '*Yeah*', '*Agree*', '*No*', '*$600*', '*Okay*', *Can you decrease?*, etc. the model initially tends to comprehend the traveler's need by requesting information; however, after a few turns, it may generate repetitive or inconsistent responses. This discrepancy arises from the training data, which comprises interactive dialogues with longer utterances, leading to confusion when handling concise inputs. Furthermore, continuous out-of-context responses may prompt the model to generate inadequate responses, given that language models inherently possess extensive knowledge in their memories.

Ethics Statement

Given the inherent bargaining nature of negotiations, developing and utilizing personalized negotiation dialogue systems must adhere to ethical conduct. Our Institutional Review Board (IRB) evaluated and approved this study. Our personalized integrative negotiation approach emphasizes traveler's flexibility, fostering a win-win outcome. Contrary to a zero-sum scenario, where one party's gain is the other's loss, the proposed negotiation dialogue system prioritizes mutual benefit. Throughout the conversation, travelers retain the autonomy to reject the deal, ensuring they are not obliged to proceed if it doesn't align with their preferences.

The ethical considerations should also be taken care of in case of the dataset. Due to the synthetic nature of the dialogues, users should approach them with sensitivity and respect and be cognizant of the potential risks of cultural appropriation or misrepresentation when generating data using language models. Given that language models are trained on web-based data, there is a potential for these models to incorporate biases that reinforce stereotypes, contribute to discrimination, or marginalize particular communities. Previous research has also reported the potential influence of synthetic data in contributing to feedback loops that intensify the occurrence of biased language generation (Taori and Hashimoto, 2023). Hence, it is imperative to engage in collaboration with linguists, language experts, and community representatives to prevent the inadvertent perpetuation of stereotypes and cultural insensitivity.

To ensure proper use, access to the created dataset will be granted upon completion and signing an agreement that the data will exclusively be used for research purposes. Human experts, regular employees of our research group engaged in the annotation, filtering/editing of data, and manual evaluations are compensated in accordance with institute policies.

Acknowledgement

The research reported in this paper is an outcome of the project titled "Conversational Agents with Negotiation and Influencing Ability", sponsored by Accenture Labs, Banglore, India.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Zishan Ahmad, Suman Saurabh, Vaishakh Menon, Asif Ekbal, Roshni Ramnani, and Anutosh Maitra. 2023. Ina: An integrative approach for enhancing negotiation strategies with reward-based dialogue agent. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2536–2549.

Nicholas Asher, Julie Hunter, Mathieu Morey, Farah Benamara, and Stergos Afantenos. 2016. Discourse structure and dialogue acts in multiparty dialogue: the stac corpus. In *10th International Conference on*

Language Resources and Evaluation (LREC 2016), pages 2721–2727.

Guy Assaker and Rob Hallak. 2013. Moderating effects of tourists' novelty-seeking tendencies on destination image, visitor satisfaction, and short-and long-term revisit intentions. *Journal of Travel Research*, 52(5):600–613.

Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.

Maria Boritchev and Maxime Amblard. 2022. A multi-party dialogue ressource in french. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 814–823.

Peter F Brown, Stephen A Della Pietra, Vincent J Della Pietra, Jennifer C Lai, and Robert L Mercer. 1992. An estimate of an upper bound for the entropy of english. *Computational Linguistics*, 18(1):31–40.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Kushal Chawla, Gale Lucas, Jonathan May, and Jonathan Gratch. 2022. Opponent modeling in negotiation dialogues by related data adaptation. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 661–674.

Kushal Chawla, Jaysa Ramirez, Rene Clever, Gale Lucas, Jonathan May, and Jonathan Gratch. 2021. Casino: A corpus of campsite negotiation dialogues for automatic negotiation systems. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3167–3185.

Yang Deng, Lizi Liao, Liang Chen, Hongru Wang, Wenqiang Lei, and Tat-Seng Chua. 2023a. Prompting and evaluating large language models for proactive dialogues: Clarification, target-guided, and non-collaboration. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10602–10621.

Yang Deng, Wenxuan Zhang, Wai Lam, See-Kiong Ng, and Tat-Seng Chua. 2023b. Plug-and-play policy planner for large language model powered dialogue agents. In *The Twelfth International Conference on Learning Representations*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *Preprint*, arXiv:1810.04805.

Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmad, and Li Deng. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 484–495.

Peyman Faratin, Carles Sierra, and Nick R Jennings. 1998. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159–182.

Yao Fu, Hao Peng, Tushar Khot, and Mirella Lapata. 2023. Improving language model negotiation with self-play and in-context learning from ai feedback. *arXiv preprint arXiv:2305.10142*.

Sebastian Gehrmann, Hendrik Strobelt, and Alexander M Rush. 2019. Gltr: Statistical detection and visualization of generated text. *arXiv preprint arXiv:1906.04043*.

Robert H Guttman and Pattie Maes. 1998. Agent-mediated integrative negotiation for retail electronic commerce. In *International Workshop on Agent-Mediated Electronic Trading*, pages 70–90. Springer.

He He, Derek Chen, Anusha Balakrishnan, and Percy Liang. 2018. Decoupling strategy and generation in negotiation dialogues. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2333–2343.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751*.

Youngsoo Jang, Jongmin Lee, and Kee-Eung Kim. 2022. Gpt-critic: Offline reinforcement learning for end-to-end task-oriented dialogue systems. In *10th International Conference on Learning Representations, ICLR 2022*. International Conference on Learning Representations, ICLR.

John F Kelley. 1984. An iterative design methodology for user-friendly natural language office information applications. *ACM Transactions on Information Systems (TOIS)*, 2(1):26–41.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Cristina Mihaela Lazăr et al. 2019. Internet–an aid for e-tourism. *Ecoforum*, 8(1):0–0.

Mike Lewis, Denis Yarats, Yann N Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning for negotiation dialogues. *arXiv preprint arXiv:1706.05125*.

Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2015. A diversity-promoting objective function for neural conversation models. *arXiv preprint arXiv:1510.03055*.

Jiwei Li, Will Monroe, Alan Ritter, Michel Galley, Jianfeng Gao, and Dan Jurafsky. 2016. Deep reinforcement learning for dialogue generation. *arXiv preprint arXiv:1606.01541*.

Yu Li, Kun Qian, Weiyan Shi, and Zhou Yu. 2020. End-to-end trainable non-collaborative dialog system. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8293–8302.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3):276–282.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.

Sangwon Park, Iis P Tussyadiah, Josef A Mazanec, and Daniel R Fesenmaier. 2010. Travel personae of american pleasure travelers: a network analysis. *Journal of travel & tourism marketing*, 27(8):797–811.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms.

Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. 2019. What makes a good conversation? how controllable attributes affect human judgments. *arXiv preprint arXiv:1902.08654*.

Weiyan Shi, Yu Li, Saurav Sahay, and Zhou Yu. 2021. Refine and imitate: Reducing repetition and inconsistency in persuasion dialogues via reinforcement learning and human demonstration. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3478–3492.

Rohan Taori and Tatsunori Hashimoto. 2023. Data feedback loops: Model-driven amplification of dataset biases. In *International Conference on Machine Learning*, pages 33883–33920. PMLR.

MN Team et al. 2023. Introducing mpt-7b: a new standard for open-source, commercially usable llms.

Leigh L Thompson, Jiunwen Wang, and Brian C Gunia. 2010. Negotiation. *Annual review of psychology*, 61:491–515.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Xiaoyang Wang, Chen Li, Jianqiao Zhao, and Dong Yu. 2021. Naturalconv: A chinese dialogue dataset towards multi-turn topic-driven conversation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14006–14014.

Bernard L Welch. 1947. The generalization of 'student's' problem when several different population varlances are involved. *Biometrika*, 34(1-2):28–35.

Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2018. Dialogue natural language inference. *arXiv preprint arXiv:1811.00671*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface's transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Qingyang Wu, Yichi Zhang, Yu Li, and Zhou Yu. 2021. Alternating recurrent dialog model with large-scale pre-trained language models. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1292–1301.

Atsuki Yamaguchi, Kosui Iwasa, and Katsuhide Fujita. 2021. Dialogue act-based breakdown detection in negotiation dialogues. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 745–757.

Runzhe Yang, Jingxiao Chen, and Karthik Narasimhan. 2020. Improving dialog systems for negotiation with personality modeling. *arXiv preprint arXiv:2010.09954*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B Dolan. 2020a. Dialogpt: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278.

Zheng Zhang, Lizi Liao, Xiaoyan Zhu, Tat-Seng Chua, Zitao Liu, Yan Huang, and Minlie Huang. 2020b. Learning goal-oriented dialogue policy with opposite agent awareness. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 122–132.

Tiancheng Zhao, Kaige Xie, and Maxine Eskenazi. 2019. Rethinking action spaces for reinforcement learning in end-to-end dialog agents with latent variable models. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1208–1218.

Yiheng Zhou, He He, Alan W Black, and Yulia Tsvetkov. 2019. A dynamic strategy coach for effective negotiation. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 367–378.

Frequently Asked Questions (FAQs)

- **What factors led to the selection of the MPT-7B model over other models, including those with more parameters for dataset creation?**

  For selecting the LLM for dataset creation, we conducted experiments using GPT-J 6B, Falcon 7B, and Alpaca 7B models for dialogue generation. After generating 50 conversations with each LLM, we found that the MPT-7B model consistently produced higher-quality conversations in larger quantities than the GPT-J 6B, Falcon 7B, and Alpaca 7B models. Also, there is no denying fact that with the regular emergence of new LLMs, there is a prevailing conviction that forthcoming iterations will exhibit better performance. However, due to the computational resource limitations, we have not used other LLMs with more parameters demonstrating an improved performance for various tasks.

- **Why did you limit your experiments to the proposed DEAL dataset?**

  We conducted experiments solely on the proposed DEAL dataset due to the absence of a dialogue dataset specifically tailored for negotiating various factors (price, aspects, amenities) within the tourism domain. To address this gap, we created the DEAL dataset and meticulously annotated it with relevant 'travel personas', intents, negotiation strategies, and negotiation phases for personalized negotiation conversations. Our preliminary analysis of the dataset highlights the diverse travel personas, intents, strategies, and phases expressed in the utterances. Therefore, we opted to utilize the newly created dataset in this study. In future endeavors, we intend to create and incorporate additional datasets by employing thorough annotations and applying our methodologies.

- **How do you anticipate the applicability of the proposed methodology across various domains?**

  We envision that the proposed methodology, based on a reinforcement learning framework, holds significant potential for application across diverse domains. By leveraging the flexibility inherent in reinforcement

learning, particularly through modifying trait-focused rewards, our approach can be tailored to address the unique challenges and objectives of negotiation within different domains like sales and job interviews, to mention a few.

- **Why was the comparison limited to only eight methods?**

  We opted to evaluate some of the most widely recognized methods in the field. While it's conceivable that other contemporary approaches may exist that we did not assess, it's improbable that such methods would also be highly effective in the context of personalized integrative negotiation conversations.

- **What measures were taken to ensure the realism and ethical integrity of the dataset utilized in the study, considering that the data was generated through prompts to Large Language Models (LLMs)?**

  A combination of careful prompt design, human oversight, content filtering, and validation procedures was employed to mitigate ethical concerns associated with the dataset created using prompting LLMs. The prompts given to LLMs were carefully crafted to prevent inappropriate or biased responses. Human annotators thoroughly verified the dataset to remove any content that could perpetuate bias or offensive language. Any questionable or problematic content instances were excluded from the dataset. Additionally, extensive testing and validation ensured the dataset's suitability for research while addressing ethical concerns.

- **Did encounter any challenges related to known issues such as reward hacking or instability arising from optimizing multiple metrics?**

  Reward hacking and instability are common challenges encountered in training RL-based models, especially when optimizing multiple metrics. To prevent reward hacking, we meticulously designed reward functions that accurately reflect the desired task objectives so that these are less susceptible to exploitation. Furthermore, we constantly monitor the training to identify instances of reward hacking and adjust the training process accordingly.

- **Also, how did the coefficients of each weight were chosen?** One of the significant challenges stems from optimizing reward weights, potentially extending both training and validation times. To mitigate this, heuristic approaches are employed to select specific combinations of reward weights (highlighted in the Limitations section). Finally, the coefficients of each weight for different rewards are chosen empirically. The results of experiments with different reward weights are presented in Table 13 in the appendix.

Appendix

This section presents complementary materials, including detailed dataset creation procedures, RL policy optimization, implementation details, and a discussion on the conversation samples to enhance the reader's comprehension of the work.

A Dataset Details

A.1 Travel Packages

The background database consists of information about various travel packages, namely Adventure Trekking Expedition, Cultural Heritage Tour, Beach Getaway, Wildlife Safari, Wellness Retreat, Historical Architecture Tour, Mountain Resort, Cruise Package, Education Tour, Eco Tour, Mountain Adventure, Cultural Exploration, Tropical Beach Retreat, Wilderness Adventure Expedition, Culinary Tour, Romantic Escape Honeymoon Package, Capture the Moments Photography Tour, Festive Splendor: A Celebration Tour, Family Adventure Expedition: Fun for Everyone and Art and Architecture Discovery: Unveiling Masterpieces.

A.2 Prompt Formulation

The proposed DEAL dataset comprises conversations focusing on negotiation on different features associated with tourism packages, including but not limited to price. To create synthetic data focusing on such aspects, we utilize a four-shot prompt (owing to the token limit of 2,048 in MPT-7B) to prompt the MPT-7B model. Each few-shot contains a task description, required context (summary), and an example utterance. For each intent, 3-shot examples are constructed manually and the 4-th shot (which is the target shot), we create a template-based prompt with various slots to be filled in depending upon the intent. An example of 4-shot prompt for the travel agent's intent 'Negotiate remove-X' is shown in the box. In this case, each shot contains the information about the amenity to be removed and final deal price after amenity removal. The value of the 'amenity-to-remove' depends on the customer's preference and the 'final-deal-price-after-removal' is calculated by fetching the cost of the removed amenity from the background database and then deducting it from the current seller price.

Prompt for agent's intent 'Negotiate remove-X'

[Task Description] You are a travel agent, you are in negotiation with the traveler for a tour package *Wildlife Safari*, which has an exciting wildlife adventure, exploring diverse habitats and encountering fascinating wildlife species. You are proposing budget for tour package as $37,750.
[Summary] The initial deal was a tour package for *Wildlife safari* with guided night safari tours, photography workshops, and binoculars rental. The price for this deal was $37,750. The traveler wants to remove an amenity binoculars rental from the package.
The traveler says "*I want to remove binoculars rental from the package.*" Then you will say "*Sure, after removing the amenity binoculars rental the cost of tour package is $34000*".

[Task Description] You are a travel agent, you are in negotiation with the traveler for a tour package *Mountain Resort*, with a tranquil mountain retreat, surrounded by breathtaking natural beauty and offering a range of thrilling outdoor activities. You are proposing budget for tour package as $32,750.
[Summary] The initial deal was a tour package for *Mountain Resort* with scenic rides, hiking tools, and snowboarding. The price for this deal was $32,750. The traveler wants to remove an amenity Hiking tools from the package.
The traveler says "*I would like to remove Hiking tools in the package.*"
Then you will say "*Okay sure, after removing the amenity Hiking tools, package cost is $23000*".

[Task Description] You are an agent, you are in negotiation with the customer for a tour package *Wildlife Safari*, which embarks on an exciting wildlife adventure, exploring diverse habitats and encountering fascinating wildlife species. You are proposing budget for tour package as $43,750.
[Summary] The initial deal was a tour package for *Wildlife safari* with walking safari guide, camera rental, and birdwatching opportunities. The price for this deal was $43,750. The traveler wants to remove an amenity camera rental from the package.
The traveler says "*I need the Camera rental amenity, what it will cost.*" Then you will say "*Okay, if we remove the amenity Camera rental, the final tour cost is $38000*".

[Task Description] You are a travel agent, you are in negotiation with the traveler for a tour package *Wellness Retreat*, which has an embarking rejuvenating journey focused on wellness, self-care, and relaxation, nurturing your mind, body, and soul. You are proposing budget for tour package as $15,000.
[Summary] The initial deal was a tour package for *Wellness Retreat* with healing therapies, relaxation areas and spa access. The price for this deal was $15000. The traveler wants to remove an amenity Relaxation Areas from the package.
The travelers says "*I don't want to keep relaxation areas in my amenities, can you please remove it. After removing Relaxation, Areas the tour cost would be $14,200.*"
Then you will say <generated_response>

A.3 Turing Test for DEAL

We compare our LLM-generated dialogues with human-generated ones to compare the quality of the conversations generated by our dataset generation approach against those generated by humans. For this, we randomly select 45 dialogues from the synthetically generated DEAL dataset along with the background information that was used to generate those dialogues using LLM. This information is then given to three human experts who have post-graduate qualifications in English linguistics and substantial experience in related tasks. These experts are randomly paired with each other to ensure variability in the written dialogues, and each pair of experts is asked to generate 15 dialogues each. Based on the provided background information, the experts are instructed to create negotiation dialogues between a travel agent and a traveler using a well-known Wizard-of-Oz approach (Kelley, 1984), where one human assumes the role of the travel agent, and the other acts as the traveler. The experts are asked to negotiate for the provided travel package and its associated aspects, amenities, and services by showing their preferences and priorities towards these different aspects, amenities, and services.

Following (Gehrmann et al., 2019), we then perform a human experiment to try and detect AI-generated content. In particular, we conduct a Turing test to identify the synthetically generated dialogues. For each background information sample, we present pairs of synthetically generated dialogue and its corresponding human-written dialogue to human evaluators and ask them to select the conversation that was synthetically generated. For a robust evaluation, every pair of dialogues is annotated by 3 human evaluators, and the majority vote is used as the final annotation. The results of the Turing test are as follows: Lose: 6.67%, Win: 71.11%, and Tie: 22.22%. Interestingly, we observe that 93% of the synthetically generated dialogues have been assessed to be as human-like as those created by actual humans. We observe the Kappa agreement ratio of 78.6%, which indicates the fair agreement among the evaluators.

A.4 Negotiation Intents, Strategies and Phases

We furnish definition and example utterances for different intents in Table 6 and Table 7, respectively. Further, we provide the definition and examples for negotiation strategies in Table 8 and negotiation phases in Table 9. A sample conversation with travel personas, intents, negotiation strategies, and negotiation phases is depicted in Table 5.

A.5 Dataset Annotation Details

This section presents the details of the annotation procedure for traveler's travel persona profiles and agent's negotiation strategies and negotiation phases. The annotation process involves the collaboration of three annotators, consisting of two Ph.D. holders in Linguistics and one with a Master's degree in Computer Science[11]. All three annotators exhibit proficiency in English, considerable expertise in labeling tasks, and a comprehensive grasp of negotiation concepts.

Travel Persona Annotation. To annotate the traveler's utterance with travel persona information, we devise the *Persona Extraction Module* (PEM). PEM inputs the traveler's utterances and list of available amenities and services and outputs a list comprising the *travel persona* information. For a given utterance, PEM first uses the Spacy part-of-speech tagging module[12] to identify the noun phrases in the utterance. It then selects the aspects associated with the target utterance by calculating the cosine similarity between the extracted noun phrases and aspects retrieved from the available amenities and services. Afterward, it identifies the sentiment associated with each selected aspect as positive, neutral, or negative by employing an open-source tool for aspect-based sentiment analysis[13]. The list of aspects with positive sentiment constitutes the '*travel persona*' profile of the target utterance. This extracted persona information is cross-verified by the annotators. We achieve an agreement ratio (McHugh, 2012) of 79.4% among these annotators, which can be considered reliable.

Negotiation Strategy and Phase Annotation. The entire annotation process for negotiation strategy and phase labels proceeds in two steps to lessen human involvement. In the first step, we randomly sample 387 dialogues from the dataset and then instruct the annotators to manually annotate the travel agent's utterances with the pertinent negotiation strategy and phase labels. Afterward, we build negotiation strategy and negotiation phase

[11] The annotators are compensated according to institutional guidelines. These annotators are different from those involved in the dataset creation.

[12] https://spacy.io/api/tagger

[13] https://huggingface.co/yangheng/deberta-v3-base-absa-v1.1

Budget	Price Range	k
Low Budget	traveler initial price <= 0.65 × agent initial price	1.2
High Budget	traveler initial price >= 0.85 × agent initial price	0.6
Moderate Budget	0.85 × agent initial price > traveler initial price > 0.65 × agent initial price	0.9

Table 4: Based on the initial price proposed by traveler and agent we get the initial price gap. If this price gap is large then we consider it as the scenario of low budget, if this price gap is low then we consider it as the scenario of high budget, otherwise we consider it to be moderate budget. In this table, we have presented the relationship between traveler's initial price and agent's initial price for all three scenarios. Also we have showed the value of concession factor k corresponding to low, high and moderate budget.

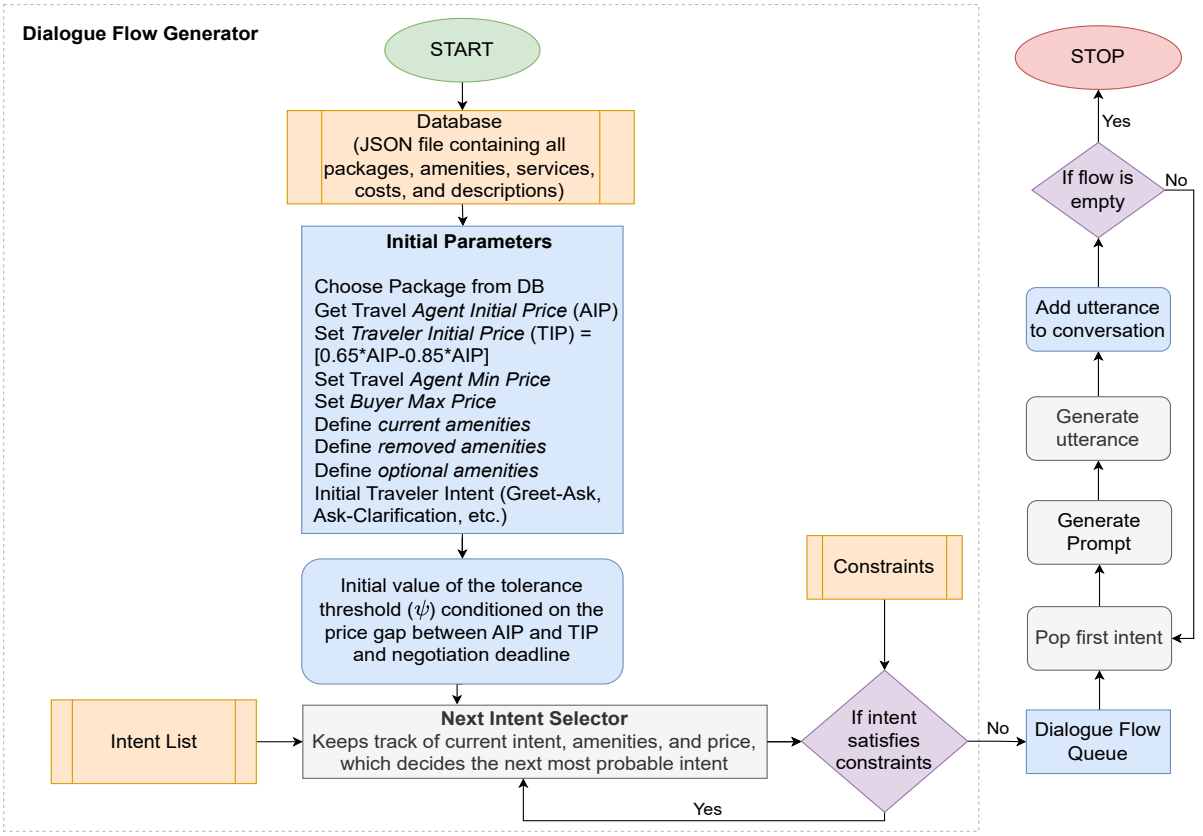

Figure 3: DEAL dataset creation process.

classifiers by fine-tuning two different pre-trained RoBERTa-large (Liu et al., 2019) models on these manually annotated samples. The utterances in the remaining dialogues are then passed through the respective classifiers to predict the negotiation strategy and negotiation phase labels. In the second step, we ask the same annotators to cross-verify and correct the predicted labels if required.

A reliable multi-rater Kappa (McHugh, 2012) agreement ratios of 78.4% and 82.6% are observed in the first step and 82.7% and 87.1% are observed in the second step for negotiation strategy and negotiation phase annotations, respectively.

A.6 Comparison with Existing Negotiation Datasets

We compare our proposed DEAL dataset with several existing dialogue datasets for negotiation along four different dimensions, *viz.* negotiation type, domain, negotiation elements/factors, and dataset size (# of dialogues and avg. # of utterances). Table 10 lists existing negotiation dialogue benchmarks along these dimensions. All these datasets either follow an integrative or distributive negotiation approach. Integrative negotiation aligns with win-win outcomes, aiming to cultivate mutual benefit. Distributive negotiation, on the other hand, typically entails win-lose scenarios, focusing on maximizing individual gains.

Speaker	Utterances
Traveler	What is included in the Education Tour package? (Ask-Clarification-Y)
Travel Agent	The Education Tour package includes accommodation, transportation, meals, educational visits, language classes, guided study sessions, academic lectures for visited sites, cultural immersion, books and study guides, and collaborative project mentorship. (Provide-Clarification-Y) (problem-solving) (problem-identification)
Traveler	I would like to enhance my learning experience, so please add Guided Study Sessions to the package. (learning experience, guided study sessions)(Negotiate-Add-X)
Travel Agent	Absolutely, adding Guided Study Sessions to the package will greatly enhance your learning experience. The final tour cost with this amenity is dollar $18900. (Negotiate-Add-X) (strategic-proposal) (problem-identification)
Traveler	I would like to remove the Cultural Immersion activity from the package as it is not something I am interested in. (learning experience, guided study sessions)(Negotiate-Remove-X)
Travel Agent	I understand that you are not interested in the Cultural Immersion activity. In that case, I would recommend the Collaborative Project mentorship as a low budget alternative. This will bring down the cost of the tour package to dollar $18900. (Negotiate-Remove-Y-Add-X) (strategic-proposal) (problem-identification)
Traveler	No, I do not want to add the Collaborative Project mentorship as an alternative. (learning experience, guided study sessions)(provide-consent)
Travel Agent	Thanks for your consent, we are not adding the Collaborative Project mentorship as an alternative. (consent-response) (strategic-proposal) (problem-identification)
Traveler	I'm sorry, but I would really like to have the Collaborative Project mentorship included in the package. Can you please add it and let me know if the budget can be adjusted accordingly? (learning experience, guided study sessions, collaborative project mentorship)(Negotiate-Price-NoChange)
Travel Agent	Adding the Collaborative Project mentorship to the package will increase the budget. After recalculating, the new budget for the package with the added amenity will be dollar $20500. Please let me know if this fits within your budget. (Negotiate-Price-NoChange) (firm-pricing) (problem-identification)
Traveler	I'm sorry, but I cannot go beyond my proposed budget of dollar $14800. (learning experience, guided study sessions, collaborative project mentorship)(Negotiate-Price-NoChange)
Travel Agent	I understand that you have a specific budget in mind. However, considering the comprehensive package we are offering which includes accommodation, transportation, meals, educational visits, language classes, and amenities such as guided study sessions, academic lectures, cultural immersion, and collaborative project mentorship, the package is priced at dollar $18840. We believe this package will provide you with a valuable educational and cultural experience. If a lower budget is a concern, we can discuss alternatives and make adjustments accordingly. (Negotiate-Price-NoChange) (firm-pricing) (problem-identification)
Traveler	Yes, I would like to proceed with the comprehensive package. (learning experience, guided study sessions, collaborative project mentorship)(provide-consent)
Travel Agent	Thanks for your consent, we will proceed with the comprehensive package. (consent-response) (strategic-proposal) (problem-identification)
Traveler	Thank you for considering my proposal. I look forward to experiencing the comprehensive package. (learning experience, guided study sessions, collaborative project mentorship)(Negotiate-Price-NoChange)
Travel Agent	Thank you for choosing our tour package! We are excited to have you as our customer. Let's close the deal and get ready for an amazing experience. (Accept) (definitive-decision-making) (decision-making)

Table 5: A sample conversation between the traveler and travel agent engaged in negotiation with persona, intent, negotiation strategy, and negotiation phase annotations.

The *STAC* (Asher et al., 2016) benchmark assesses player communication in the game of Catan, where players collect resources like wood, wheat, and sheep in an integrative manner to build settlements, roads, and cities. *DealorNoDeal* (Lewis et al., 2017) is another benchmark dataset for integrative negotiation dialogues in which each participant is given their own set of items, each assigned a value, and both aim to maximize their overall score through negotiation. *CraigslistBargain* (He et al., 2018) and *NegoCoach* (Zhou et al., 2019) datasets consist of distributive dialogues based on a realistic item price bargaining scenario. *Anti-Scam* (Li et al., 2020) focuses on distributive negotiation in online customer service domain where participants seek to safeguard themselves by identifying potential attackers attempting to steal sensitive personal information.

CaSiNo (**?**) dataset involves campsite scenarios where campsite neighbors integratively negotiate for additional food, water, and firewood packages, and both parties have different priorities over different items. *JobInterview* (**?**) dataset involves integrative recruiter-applicant dialogues negotiating over salary, day off, position, company, and workplace. *IND* (Ahmad et al., 2023) consists of integrative negotiation dialogues where the deal is modeled as a bundle of electronic products and related accessories.

Our proposed DEAL dataset differs from the existing negotiation dialogue datasets along all the provided dimensions. DEAL stands out due to its focus on tourism-related aspects and amenities, encompassing topics, such as price, destination, entertainment, transportation, hotels, and other relevant features. With 1,291 dialogues and an average dialogue length of 14.66, DEAL offers a unique perspective on integrative negotiations within the tourism domain. Besides, the proposed dataset provides exhaustive annotation for travel persona,

Intent	Definition
Greet	Initiating the conversation politely and setting a positive tone for the dialogue.
Ask	Asking for information about a particular aspect, amenity, travel package, or ongoing negotiation.
Inform	Providing detailed information about the packages or services involved in the negotiation.
Ask price	Inquiring or clarifying a package or service's cost or pricing terms.
Tell price	Communicating a package or service's proposed price or cost.
Ask clarification-Y	Acquiring additional information or clarification on specific aspects of the negotiation or the overall deal according to the current negotiation state.
Provide clarification-Y	Clarifying specific facets of the negotiation or the comprehensive deal in alignment with the present state of the negotiation process.
Negotiate price increase	Involves the travel agent's endeavor to raise the package or service's pricing terms within the ongoing negotiation.
Negotiate price decrease	Involves the travel agent's aim to reduce the pricing terms of a package or service during negotiation.
Negotiate price no-change	Involves the agent's inclination to propose or assert that the price of a package or service remains unchanged by emphasizing the value and equity of the current deal.
Negotiate add-X	Proposing or suggesting adding a specific aspect or service during negotiation.
Negotiate remove-X	Proposing or suggesting the removal of a specific aspect or service from the deal, which may or may not lead to a decrease in the deal's price.
Provide consent	Agreeing or approving a proposal, condition, or term during the negotiation.
Consent response	Acknowledging and responding to the provided consent. It signifies the agent's understanding of the traveler's agreement and readiness to proceed accordingly.
Accept	Acceptance of the proposal, offer, deal, or situation attained during negotiation.
Reject	Rejection of the proposal, offer, deal, or situation attained during negotiation.
Acknowledge acceptance	Acknowledging the acceptance of the deal.

Table 6: Intents definition. The intents may manifest independently or in conjunction with others, for example, '*Greet-Ask*', '*Negotiate remove-X_1 add-X_2*. Examples for each intent are provided in Table 7 of the appendix.

Intent	Example
Greet-Ask	*Hello, I am interested in purchasing the Mountain Resort package. Could you provide me with details about the package and its price?*
Greet-Inform	*Hello, the Mountain Resort package offers luxurious accommodation, convenient transportation, and delicious meals. In addition, you can enjoy amenities such as scenic rides, snowboarding, mountain biking, rock climbing, indoor swimming pool, and gourmet dining. The price of this package is $19750.*
Ask price	*Could you please provide me with the total cost of the tour package?*
Tell price	*The package will cost you $13200?*
Ask clarification-Y	*Please tell me what type of room will be there in Mountain Resort*
Provide clarification-Y	*Sure sir, this package is offering you rooms which are called Deluxe Double or Twin Room, Superior Double or Twin Room, Superior Double or Twin Room, Standard Double or Twin Room, Superior Double or Twin Room, Standard Double or Twin Room, Deluxe Triple Room, Superior Triple Room, Superior Triple Room, Standard Triple Room, Superior Triple Room, Superior Double or Twin Room and Deluxe Triple Room.*
Negotiate price increase	*Thank you for your interest in our tour package and appreciating the detailed information. I understand that your budget is $16284, and I would like to accommodate your request as much as possible. After considering your budget, I can offer you a slightly adjusted package for $17821. Let me know your thoughts on this offer.*
Negotiate price decrease	*I appreciate the detailed information about the tour package. However, my budget for this educational tour is $16284. Is there any possibility of adjusting the package to fit within my budget?*
Negotiate price no-change	*As this package is all in one place of adventures, you really need to give me a chance to provide you more exciting options for your budget.*
Negotiate add-X	*Sir I would like to recommend the amenity to go for local guides tour which will be around $16900.*
Negotiate remove-X	*I don't require any food pairing during the tour, so please remove that amenity from the package.*
Negotiate remove-X1 add-X2	*I understand that you don't require any food pairing during the tour. As an alternative, I can offer you a Farm-to-table Experience which will lower down the price of the tour to $12041.*
Provide consent	*Yes, I would like to add the Farm-to-table Experience.*
Consent response	*Thanks for your consent, we are adding the Farm-to-table Experience as a service.*
Accept	*I'm really happy with the package and the proposed price, so let's go ahead and finalize the deal. Thanks!*
Reject	*Sorry Sir, we cannot provide you this package in such lower cost.*
Acknowledge acceptance	*I'm happy to inform you that the tour package deal has been accepted. Thank you for choosing our services!*

Table 7: The example utterances of different intents in the proposed DEAL dataset.

Negotiation Strategy	Definition	Example
Problem identification	It refers to understanding travelers' key issues to clarify their concerns, needs, and expectations to lay the foundation for achieving mutually agreeable solutions.	*Hello, the Mountain Resort package offers luxurious accommodation, convenient transportation, and delicious meals. In addition, you can enjoy amenities such as scenic rides, snowboarding, mountain biking, rock climbing, indoor swimming pool, and gourmet dining. The price of this package is $19750.*
Strategic proposal	It refers to presenting proposals to achieve specific objectives, often with a focus on gaining advantages or maximizing outcomes for the travel agent.	*Absolutely, adding Guided Study Sessions to the package will greatly enhance your learning experience. The final tour cost with this amenity is $18900.*
Firm pricing	It refers to presenting non-negotiable, fixed terms or prices. This strategy is characterized by minimal flexibility and is employed when the travel agent believes the presented price accurately reflects the true value of the offering.	*I understand that your budget is limited, but unfortunately, we cannot adjust the package to fit within your budget. However, I can offer you a discounted rate of $19025 for the luxurious accommodations at the Mountain Resort.*
Definitive decision-making	It emphasizes making clear and final decisions during negotiation. It involves a more authoritative and decisive approach to reaching agreements.	*Thank you for considering our tour package deal. We are happy to accept and close the deal.*
Collaborative proposal	It is characterized by the joint effort of the travel agent and the traveler to suggest, discuss, and potentially modify various aspects of a travel package. The travel agent emphasizes the active involvement of the traveler in shaping the proposed terms of the travel arrangement.	*I understand that you don't require any food pairing during the tour. As an alternative, I can offer you a Farm-to-table Experience which will lower down the price of the tour to $12041.*
Flexible pricing	It allows for price adjustments and modifications during the negotiation. This strategy is more adaptable and open to accommodating price adjustments, providing a balanced approach to reaching mutually agreeable pricing.	*Considering your request about the tour package, we can have concession for you, it will cost you $18300*
Co-operative decision-making	It involves making decisions through cooperation and consensus-building to ensure that decisions are acceptable to both the traveler as well as the travel agent.	*Thanks for your patience, let's close the deal*
No strategy	It is designated to the utterances that do not employ any negotiation strategy.	*I'm happy to inform you that the tour package deal has been accepted. Thank you for choosing our services!*

Table 8: The definition and example utterances of different negotiation strategies in the proposed DEAL dataset.

Negotiation Phase	Definition	Example
Relational positioning	This phase concentrates on the preparation and strategic planning for negotiations, characterized by the concerted efforts of the travel agent and the traveler to recognize a dispute, articulate an agenda, underscore points of divergence, and assume specific positions.	*Hello, the Mountain Resort package offers luxurious accommodation, convenient transportation, and delicious meals. In addition, you can enjoy amenities such as scenic rides, snowboarding, mountain biking, rock climbing, indoor swimming pool, and gourmet dining. The price of this package is $19750.*
Problem-solving	This phase prioritizes seeking a resolution for the dispute and is characterized by strategic maneuvers, extensive interpersonal interaction, and a gradual shift towards achieving a mutually acceptable agreement.	*I understand that you are not interested in the Cultural Immersion activity. In that case, I would recommend the Collaborative Project mentorship as a low budget alternative. This will bring down the cost of the tour package to $18900.*
Offer generation	This phase centers on finding and proposing offers that meet the traveler's requirements.	*I understand that your budget is limited, but unfortunately, we cannot adjust the package to fit within your budget. However, I can offer you a discounted rate of $19025 for the luxurious accommodations at the Mountain Resort.*
Decision making	This phase focuses on meticulously considering details to finalize and implement the agreement (close the deal).	*Thank you for considering our tour package deal. We are happy to accept and close the deal.*

Table 9: The definition and example utterances of different negotiation phases in the proposed DEAL dataset.

novel intents, negotiation strategies, and negotiation phases.

B Training and RL Policy Optimization with PPO

Warm-start using supervised learning. We use the pre-trained weights of DialoGPT-small (Zhang

Dataset	Negotiation Type	Domain	Negotiation Elements/Factors	# Dialogues	Avg. # of utterances
STAC (Asher et al., 2016)	Integrative	Strategy Games	Wood, wheat, and sheep	1,081	8.5
DealorNoDeal (Lewis et al., 2017)	Integrative	Item Assignment	Books, hats, and footballs	5,808	6.6
CraigslistBargain (He et al., 2018)	Distributive	E-commerce	Product Price	6,682	9.2
NegoCoach (Zhou et al., 2019)	Distributive	Product Bargaining	Price	300	-
Anti-Scam (Li et al., 2020)	Distributive	E-commerce	Privacy Protection	220	12.45
CaSiNo (Chawla et al., 2021)	Integrative	Item Assignment	Food, water and firewood	1,030	11.6
JobInterview (Yamaguchi et al., 2021)	Integrative	Job Interview	Salary, day off, position, company, and workplace	2,639	12.7
IND (Ahmad et al., 2023)	Integrative	E-commerce	Price, Electronic products (air Conditioning, television, refrigerator, oven, washing machine, printer, smartphone, laptop, tablet, and camera) and related accessories	4,163	13.79
DEAL ***(Proposed)***	Integrative	Tourism	Price, Destination, Entertainment, Transportation, Hotels, and other tourism-related aspects and amenities.	1,291	14.66

Table 10: Comparison of the proposed DEAL dataset with the existing negotiation dialogue datasets. All the benchmarks include two-party negotiation dialogues.

et al., 2020a) for initializing PSFT model. Next, we use a warm-start strategy using supervised learning on the proposed DEAL dataset, following previous work in reinforcement learning for dialogue generation (Li et al., 2016). In particular, we use DEAL dataset to fine-tune our DialoGPT-initialized PSFT model.

PPO optimization. We reinforce the yielded reward to optimize the current state using PPO. The action selection of TRIP NEGOTIATOR can be articulated by the RL-based parameterized control policy, which is a probability mapping function denoted as $P(\Theta)$. For a given dialogue context (a state), $P(\Theta)$ produces an utterance r consisting of V tokens (an action):

$$P_\Theta(r_{1:m}|x) = \prod_{v=0}^{V} P_\Theta(r_v|y_{<v}, x) \quad (4)$$

We initialize P_Θ with π_{PSFT} during RL fine-tuning. Policy updates occur at each step using the PPO-loss, which effectively minimizes variance compared to the old policy. The policy optimization is structured in three steps: (i) gradient ascent on the loss function $\mathcal{L}(\Theta)$ to maximize expected rewards (Equation 5), (ii) replacing the log term with an importance sampling term to control large deviations and utilizing clipping to prevent catastrophic forgetting, ensuring that the updates do not have excessively large magnitudes (Equation 6), and (iii) updating parameters using Equation 7.

$$\nabla_\Theta \mathcal{L}(\Theta) = \mathbb{E}_{r_\Theta}[\nabla_\Theta \log P_\Theta(r) \hat{A}_r] \quad (5)$$

$$\mathcal{L}_{\text{CLIP}}(\Theta) = \hat{\mathbb{E}}[\text{MIN}(pr_r(\Theta)\hat{A}_r, \text{CLIP}(pr_y(\Theta), 1-\varepsilon, 1+\varepsilon)\hat{A}_r)] \quad (6)$$

$$\Theta_{k+1} = \underset{\Theta}{\text{argmax}} \underset{s,a\sim P_{\theta_k}}{\mathbb{E}} [\mathcal{L}^{\text{CLIP}}] \quad (7)$$

Here, the probability ratio between the new and old policies is calculated as $pr_r(\Theta) = P_\Theta^{new}/P_\Theta^{old}$. ε and $\hat{A}_r$ signify the clipping range and normalized rewards (estimated advantage), respectively.

C Experiments

In this section, we provide a description of baselines and a detailed experimental setup.

C.1 Baseline Description

1. **GPT-2** (Radford et al., 2019). It is a language model based on the Transformer architecture and pre-trained on Reddit dialogues. This baseline model is implemented by fine-tuning the pre-trained GPT-2-small model in a supervised learning setting on the proposed DEAL dataset.

2. **DialoGPT** (Zhang et al., 2020a). The DialoGPT was trained using 147 million Reddit chats and is based on the OpenAI GPT-2 architecture. This baseline model is implemented by fine-tuning the DialoGPT-small model in a supervised learning setting on the proposed DEAL dataset.

3. **ARDM** (Wu et al., 2021). ARDM uses two GPT2-small models to model the traveler and the travel agent separately, and jointly train them in a supervised fashion to better capture different speakers' language styles.

4. **PersRFI** (Shi et al., 2021). PersRFI is a reinforcement learning framework to improve

dialogue response generation in persuasion dialogues. The framework aims to reduce repetition and inconsistency in generated responses by incorporating human demonstration and task relevance rewards. We have used this approach for response generation in negotiation dialogues with the same rewards. We leverage GPT-2-small as the language model.

5. **GPT-Critic** (Jang et al., 2022). GPT-Critic improves LM through the cloning of critic-guided self-generated sentences during fine-tuning. It is built upon the GPT-2-small language model.

6. **INA** (Ahmad et al., 2023). Integrative Negotiation Agent (INA) leverages a reward-based dialogue system. The rewards used to train the INA are specifically tailored for the negotiation task. These rewards encourage the INA to learn strategies that can dynamically adjust prices and negotiate the inclusion or exclusion of items in a bundle deal, thereby effectively engaging in integrative negotiations. Following the original work, we implement INA using GPT-2-small as the language model.

7. **ProCoT** (Deng et al., 2023a). ProCoT involves prompting the LLM-based dialogue system to generate a chain-of-thought descriptive analysis for planning the strategy for the next turn. This method originally proposed for negotiation dialogues uses a prompt[14] that we directly employ in our experiments for a fair comparison. Also, following the work, we utilize a static version of ChatGPT, GPT-3.5-TURBO-0301, and set the temperature to 0 to produce deterministic outputs with the same inputs and set the maximum number of new tokens to 128 for the generation. In addition, we employ an open-source LLM, i.e., LLaMA-7b-chat (Touvron et al., 2023), for the evaluation. All the experiments are performed with one-shot demonstrations.

8. **PSFT**. This baseline model is built by fine-tuning the DialoGPT-small in a supervised setting with the traveler's travel persona profile.

C.2 Implementation Details

All the implementation is done using the PyTorch[15] framework. The transformer-based models from HuggingFace (Wolf et al., 2019) are utilized in all the experiments.

For BERT-based NLI model, the initialization of the BERT model involved employing the weights of 'bert-base-uncased'. The training is done with an initial learning rate of $1e^{-4}$, employing a linear schedule and incorporating a warmup technique (Vaswani et al., 2017), utilizing the Adam optimizer (Kingma and Ba, 2014). The training of the NLI model is done with batch size of 8.

All the classifiers are built considering a dialogue context length of 4, which is chosen empirically. We experiment with a context length of $\{0, 2, 4, 6\}$ for all the classifiers. Figure 4 presents the results of classifiers with varying context lengths. The classifiers are trained for ten epochs with a batch size of 32. The intent-phase models are trained for 8 epochs with a batch size of 16.

The language models GPT-2-small (Radford et al., 2019), DialoGPT-small (Zhang et al., 2020a), and ARDM (Wu et al., 2021) are trained with a cross-entropy loss in the supervised learning setting. AdamW optimizer (Kingma and Ba, 2014) is employed with a learning rate of $\alpha = 2e^{-05}, \varepsilon = 0.2$, and seed value of 10.

TRIP NEGOTIATOR is trained in an RL framework using a fine-tuned PSFT model on the DEAL dataset. Experiments are conducted with distinct values of $n = \{2, 3, 4, 5, 8\}$ (n denotes the number of candidate responses), and the optimal value of $n = 3$ is determined based on the PPL score. Nucleus sampling (Holtzman et al., 2019) is employed for decoding the generated utterances. The decoding process utilizes a temperature $\mathcal{T} = 0.8$ and a probability $p = 0.9$. The final values of coefficients are set empirically as: $w_{pa} = 0.3, w_{nsc} = 0.2, w_{ipa} = 0.2, w_{rf} = 0.1, w_{dc} = 0.1, w_d = 0.1, \eta = 2$ and $\delta = 2$. For training the TRIP NEGOTIATOR in an RL framework, we set the batch size to 8, seed value to 10, human reward to 10, max candidate length to 50, clip ratio to 0.2, discount factor to 0.95, number of steps to 32,000, steps per update to 640. AdamW optimizer (Kingma and Ba, 2014) with a learning rate of $\alpha = 2e^{-05}$, $\varepsilon = 0.2$ is employed. Finally, the model is trained for 20 epochs.

[14] https://github.com/dengyang17/LLM-Proactive

[15] https://pytorch.org/

C.3 Classification Models

To build the intent classification model (ICM), negotiation strategy classification model (NSCM), and negotiation phase classification model (NPCM), we fine-tune the pre-trained BERT-large (Devlin et al., 2018), RoBERTa-large (Liu et al., 2019), and DistilBERT-base (Sanh et al., 2019) models on the proposed DEAL dataset with the respective labels. Table 11 shows the results of intent, negotiation strategy, and negotiation phase classifiers. The results show that the classifiers built on the RoBERTa-large model outperform those built using BERT-large and DistilBERT-base for Weighted Accuracy (W-Acc) and Macro-F1 scores.

	BERT-large		RoBERT-large		DistilBERT-base	
Classifier	**W-Acc**	**Macro F1**	**W-Acc**	**Macro-F1**	**W-Acc**	**Macro F1**
ICM	82.46	71.80	86.37	78.81	83.74	72.29
NSCM	89.32	80.26	91.75	81.34	90.01	80.26
NPCM	91.78	82.63	93.24	84.02	92.10	81.24

Table 11: Evaluation results of different classifiers. All values are in %.

C.4 Intent-phase Generation Models

To build the intent-phase generation model (IPGM), we utilize the pre-trained GPT-2-medium (Radford et al., 2019) and DialoGPT-medium (Zhang et al., 2020a) model. Table 12 presents the results of the intent-phase generation models in terms of PPL, BLEU-4, METEOR, and BERTScore-F1. It is evident from the results that the intent-phase model built on the DialoGPT-medium outperforms the model built upon GPT-2-medium.

Models	**PPL**	**BLEU-4**	**METEOR**	**BERTScore-F1**
GPT-based intent-phase model	4.32	0.26	0.55	0.78
DialoGPT-based intent-phase model	3.71	0.29	0.63	0.82

Table 12: Evaluation results of intent-phase models.

C.5 Evaluation Metrics Details

C.5.1 Automatic Evaluation Details

For evaluating the baselines and proposed model, we use trait-focused metrics and language-focused metrics. Trait-focused metrics include Persona Alignment (PA), Negotiation Strategy Consistency (NSC), and Negotiation Phase Agreement (NPA). PA measures how well the generated responses align with the ground-truth travel persona of the travelers, NSC measures the consistency of the generated responses' negotiation strategy with the ground-truth responses' negotiation strategy. NPA checks whether the generated responses drive the negotiation towards the phase that adheres to the ground-truth responses' negotiation phase.

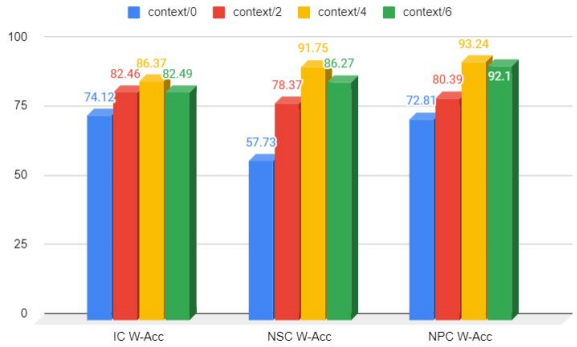

Figure 4: Effect of the context length on the performance of classifiers. IC W-Acc, NSC W-Acc, and NPC W-Acc denote the weighted accuracy scores of intent, negotiation strategy, and negotiation phase classification models.

Language-focused metrics include Perplexity (PPL) (Brown et al., 1992), traditional word-overlap-based metrics - BLEU (Papineni et al., 2002) and METEOR (Banerjee and Lavie, 2005), and and embedding-based metric - BertScore-F1 (BS-F1) (Zhang et al., 2019). PPL evaluates how well a model predicts a response. Word-overlap-based metrics like BLEU4, and METEOR compute the overlap between the ground-truth response and the model's generated response. Embedding-based metrics such as BERTScore F1 (BS-F1)[16] align the generated response and the ground-truth response in latent semantic space to assess the semantic similarity between the gold response and the model's generated response. BS-F1 computes word similarity using contextual embeddings from the BERT (Devlin et al., 2019) model.

C.5.2 Human Evaluation Details

To assess the performance of the baselines and the proposed TRIP NEGOTIATOR model from humans' perspective, we conduct human evaluation using trait-focused metrics and language-focused metrics. Trait-focused metrics include the Win Times (WT), Deal Rate (DR), and Pareto Optimality (PO) (Lewis et al., 2017) along with PA, NSC, and NPA metrics. WT represents the percentage of wins by the model. DR measures the percentage of dialogues that end up with an agreed-upon negotiation decision. PO

[16]BERTScore: https://huggingface.co/spaces/evaluate-metric/bertscore

denotes the percentage of pareto optimal solutions (a solution is Pareto optimal if neither the traveler's nor agent's score can be improved without lowering the other's score) for agreed deals. PA measures the ability of the model in terms of persona alignment, i.e., how well the model aligns with the traveler's travel persona. NSC and NPA measure the negotiation strategy correctness and negotiation phase correctness, respectively. Language-focused metrics include Fluency (F), Contextual Coherence (CC), and Engagingness (E) to evaluate the language quality of generated responses. F assesses whether the generated response is fluent (grammatically correct) and natural. CC checks whether the response is consistent with the conversational history. E measures the degree to which dialogue is engaging, compelling, and capable of maintaining the focus of the travelers.

The human evaluation is done with the help of four human evaluators[17], two evaluators with Ph.D. degrees in Linguistics, and two with post-graduate degrees in Computer Science. All the evaluators possess sufficient experience in similar tasks. Before starting the conversation with the system, the evaluators were briefed about the travel packages and their associated aspects, amenities, and service descriptions (out of the 20 travel packages that are used in the dataset preparation). Since, in the present work, we define the travel persona of the travelers in terms of their likes and dislikes towards different aspects, amenities, and/or services of a tourism package, we asked the evaluators to interact with the proposed system based on a particular travel package description and show their preferences (likes/dislikes) for various specific aspects, amenities, or services. Each evaluator was instructed to engage with the system 20 times, ensuring that each interaction involved a distinct travel package and its associated attributes with a different set of responses. The interactions were required to span at least 12 conversational turns, with multiple sentences permitted in each turn. Consequently, each evaluator contributed to 20 interactions, resulting in a total of 80 human-evaluated dialogues. Afterward, we instruct the evaluators to rate each dialogue interaction for PA, NSC, NPA, F, CC, and E on the provided scale of 1-5[18]. The higher the scores, the better the performance of the system[19].

WEIGHT OPTIMIZATION						
w_{pa}	w_{nsc}	w_{ipa}	w_{rf}	w_{dc}	w_d	**PPL**
0.1	0.2	0.15	0.25	0.05	0.25	2.7496
0.2	0.3	0.1	0.05	0.1	0.25	2.683
0.15	0.1	0.2	0.05	0.25	0.25	2.6219
0.1	0.15	0.25	0.2	0.1	0.2	2.778
0.25	0.1	0.1	0.25	0.1	0.2	2.8957
0.2	0.05	0.1	0.3	0.15	0.2	2.9321
0.1	0.25	0.05	0.2	0.15	0.25	2.746
0.3	0.1	0.2	0.05	0.1	0.25	2.734
0.2	0.1	0.2	0.15	0.1	0.25	2.634
0	0.4	0.3	0.1	0.1	0.1	2.511
0.4	0	0.3	0.1	0.1	0.1	2.4632
0.3	0.4	0	0.1	0.1	0.1	2.482
0	0	0	0.4	0.3	0.3	2.576
0.4	0.3	0.3	0	0	0	2.462
0.3	**0.2**	**0.2**	**0.1**	**0.1**	**0.1**	**2.1329**

Table 13: Weight optimization using different reward weight combinations.

C.6 Computational Resource Specifications

The experimental setup includes the following device configurations: (a) GPU: A100-PCIE-40GB, (b) CUDA Support: CUDA 11.x (or later.) (c) GPU clocks: Base: 765 MHz, Boost: 1410 MHz, (d) Memory Size: 40 GB, (e) Memory Type: HBM2, (f) Memory clock: 1215 MHz, (g) Bus Width: 5120 bits, (h) Total board power: 250 W.

C.7 Weight Optimization

To ascertain the optimal combination of weights for the reward function, we conducted experiments with different weight combinations of $w_{pa}, w_{nsc}, w_{ipa}, w_{rf}, w_{dc}$, and w_d. These weights were validated using a 10% held-out DEAL dataset, and the combination that resulted in the highest perplexity score was selected for training TRIP NEGOTIATOR. The combination yielding the highest perplexity score was subsequently chosen for training the TRIP NEGOTIATOR model. The weights considered for optimization utilizing the introduced

[17]Evaluators are distinct from those involved in the dataset creation process, and are compensated according to the institute norms.

[18]***PA, NSC, NPA*** - 1: inconsistent, 2: slightly_inconsistent, 3:irrelevant, 4: mostly consistent, 5 consistent; ***F*** - 5: Flawless, 4: Good, 3: Non-native, 2: Disfluent and 1: Incomprehensible; ***CC***- 5: All, 4: Most, 3: Much, 2: Little, and 1: Not at all; ***E***- 5: Excellent, 4: Very Good, 3: Good, 2: Bad, and 1: Not at all.

[19]The evaluations demonstrate an inter-evaluator agreement ratio of 76.34%.

DEAL dataset are outlined in Table 13. This table demonstrates that incorporating all rewards results in a lower perplexity score. Furthermore, eliminating any individual reward increases the perplexity score, underscoring each reward's significance within the model.

D Discussion

This section provides an in-depth analysis of the effectiveness and limitations of the proposed negotiation dialogue generation model, TRIP NEGOTIATOR.

D.1 Qualitative Analysis

To further assess the quality of the responses generated by the TRIP NEGOTIATOR, we present a qualitative analysis in Table 14. From the responses generated by the TRIP NEGOTIATOR, we notice that the generated responses are aligned with the travelers' personas. For instance, in Example I., it is evident that our proposed system understands the traveler's inclination to avoid the amenity '*Traditional Performances*' and effectively suggests pertinent alternatives that align with the traveler's preferences and budget constraints. This also reveals the TRIP NEGOTIATOR's efficacy in negotiating the package price. Example II. presents another scenario where the TRIP NEGOTIATOR demonstrates its ability to negotiate based on price. Besides, we notice that the negotiation strategy of the generated response is in accordance with the intended strategy. Further, we infer that the generated responses satisfy the traveler's intent and direct the conversation to the intended phase. We also observe that the grammar, semantics, and contextual relevance are well-established in the generated responses compared to the gold utterance. These analyses suggest the effectiveness of the different rewards in reward computation.

A comprehensive sample conversation illustrating actual travelers' interaction with the proposed system, TRIP NEGOTIATOR, is provided in Table 16, Table 17, and Table 18. The conversation in Table 16 depicts that the model proposes offers aligned with the traveler's persona. The conversations in Table 17 and Table 18 both pertain to the negotiation of the *Beach Resort* package. The difference between them lies in the initially proposed price by the traveler. Here, we can observe the change in the negotiation strategy adopted by the model based on the price gap. In the case of high price gap (Table 18), when the traveler's intent is *Negotiate price decrease*, the model adapts the strategy of *flexible pricing* and keep the conversation in the phase of *offer generation*. On the other hand, when the price gap is low (Table 17), and the traveler's intent is *Negotiate price decrease*, the model adapts the strategy of *firm pricing* and *collaborative proposal* and shift between the phases of *problem identification* and *offer generation*. This shows that the model successfully adopts the negotiation strategy according to the travelers' needs.

From these conversations, we can discern that the travelers give a high score to the TRIP NEGOTIATOR for various human evaluation metrics, namely, PA, NSC, NPA, F, C, and E. This indicates that the proposed system can effectively incorporate persona information and employ appropriate negotiation strategies in the generated responses. Moreover, it demonstrates the ability to adapt the negotiation phase in the responses according to the underlying intent of the traveler. These sample conversations also highlight the TRIP NEGOTIATOR's capability to produce fluent, coherent, and engaging responses.

Moreover, we present an example of traveler's interaction with our proposed system TRIP-NEGOTIATOR and the LLM's prompting-based baseline model, ProCoT+ChatGPT in Table 19. It is evident from the conversation that the traveler gives better scores for PA, NSC, NPA, F, C, and E to interaction with the proposed TRIP-NEGOTIATOR than ProCoT+ChatGPT. From the conversation, it is evident that the proposed system more effectively incorporates the traveler's preferences and engages in more strategic negotiation. For instance, for the traveler's utterance, "*It looks like there are a lot of options. However, I don't see my preferred room type in this package.*", the ProCoT+ChatGPT fails to understand the traveler's requirement and preference for a specific room type and instead offers the package at $19,100. The proposed TRIP-NEGOTIATOR, on the other hand, comprehends the traveler's need and preference for the room type and effectively accommodates this to generate a personalized response.

Further, ProCoT+ChatGPT fails to negotiate effectively. For the traveler's utterance, "*Hmm. Could you go down to $13,500?*", the traveler proposes a reduction to $13,500, and the model agrees to the reduced price of $13,500 without further negotiation. This model provides less emphasis on the value of the package and amenities

during the negotiation process. On the contrary, the proposed TRIP-NEGOTIATOR counters with a slightly reduced price of $18,500, making a deal at $18,500. The system proves efficacy in making concessions while retaining a high value for the package. Overall, these conversation illustrates that the proposed TRIP-NEGOTIATOR demonstrates a more preference-driven and strategic approach during negotiation than ProCoT+ChatGPT, resulting in effective negotiation outcomes.

D.2 Error Analysis

Though the proposed system performs fairly well, we observe certain limitations in the responses generated by TRIP NEGOTIATOR. Table 15 presents a few samples of incorrect or inconsistent responses generated by the TRIP NEGOTIATOR. For instance, in Example I., the TRIP NEGOTIATOR fails to address the traveler's concern about the package cost being beyond their budget and, thus, proposes a solution to reduce the cost while including the desired amenities. Here, the agent generates repetitive response and incorrectly suggests adding an amenity, '*Birdwatching Opportunity*'. The model occasionally generates such inconsistent responses. This may be due to the fact that our formulated reward function only considers the previous utterance for computing the dialogue coherence reward. This inconsistency may be improved by formulating a better dialogue coherence reward. In Example II., the model stumbles to understand the partial utterance that is not grammatically complete and again asks "*Would you like to add the Hot Springs to the package?*". This could be attributed to the fact that our designed reward function cannot cover the significant aspects of commonsense knowledge that can probably map "can you reduce?" to lowering the price proposed by the travel agent. Integrating pertinent commonsense knowledge may facilitate generating responses that are more reflective of real-life scenarios.

E Remark: Significance in the Era of Evolving Large Language Models

The present work, though situated within the era of GPT-2-based models, it maintains significant relevance amidst the rapidly evolving landscape of large language models (LLMs), which includes the emergence of newer architectures like LLaMA (Touvron et al., 2023), GPT-4 (Achiam et al., 2023), and many more. The core innovation of TRIP NEGOTIATOR is the reinforcement learning-based approach for personalized integrative negotiation task within the tourism domain. This approach addresses the challenge of effectively incorporating the travelers' travel persona information and necessities for the personalized integrative negotiation task using task-focused rewards. As newer LLMs like LLaMA, GPT-4, etc., continue to push the boundaries of language understanding and generation, they inherently encounter similar challenges (which is how to generate personalized integrative responses effectively). Therefore, our approach may offer a valuable technique that can be adapted and extended to these newer architectures with minimal or no modifications.

By leveraging the novel task-focused and language-focused rewards, the proposed TRIP NEGOTIATOR system effectively generates persona-aware responses adhering to the pertinent negotiation and negotiation phase aligned with the traveler's intent. Thus, while our experiments and immediate results are contextualized within the GPT-2-based architecture, the implications and potential applications of TRIP NEGOTIATOR extend far beyond. It represents a significant step forward in the ongoing evolution of language model development, emphasizing its long-term relevance in the area. Our findings not only outline a systematic approach for improving persona-aware integrative response generation in negotiation dialogues, regardless of the underlying model's architecture but also provide a road map for future research and advancements in AI and NLP, particularly in improving the flexibility and effectiveness of language models in complex, practical contexts.

F Socio-economic Implications and Reproducibility.

The present work supports tourism businesses and ongoing research endeavors by utilizing advancements in AI-based dialogue systems for negotiation. Such advancements in the tourism domain are likely to bring a high social and economic impact. Socially, these systems can enhance traveler's experience by tailoring interactions to individual preferences, providing personalized recommendations, and fostering a sense of empowerment. On the economic front, these systems can improve negotiation efficiency, benefiting both travelers and tourism businesses by saving time and resources, enhancing deal outcomes, improving traveler satis-

I.	**Context**	**traveler:** *I am interested in learning more about the tribal history, can you please include a Museum tour in the package?* **Agent:** *Certainly, adding the amenity of Museum Tour will enhance your experience. The final cost of the tour package would be $18900.* **Traveler:** *I would like to remove the Traditional performances from the package as I am not interested in them.* (Negotiate-Remove-X)
	Ground-truth	**Agent:** *It's fine if you want to remove the service Traditional performances, but as an option to this I would recommend Cultural Shopping Excursion which make this tour cost only $15800* (strategic proposal) (problem identification)
	TRIP NEGOTIATOR	**Agent:** *I understand that you want to remove Traditional performances as an amenity. If you have a budget of $0, we can recommend a lower-cost alternative which is a Cultural Immersion experience. This will bring down the price of the Tour to $18100.* (strategic proposal) (problem identification)
II.	**Context**	**traveler:** *No, I do not want to add Drone Photography Assistance as an amenity.* **Agent:** *Thanks for your consent, we are removing Drone Photography Assistance as an amenity.* **Traveler:** *Thank you for removing the Drone Photography Assistance as an amenity. However, I am still unable to go above my proposed budget of $11453. Is there any other way we can adjust the package to meet my budget?* (Negotiate-Price-NoChange)
	Ground-truth	**Agent:** *Thank you for your consideration, let's finalize the deal.* (definitive decision) (decision making)
	TRIP NEGOTIATOR	**Agent:** *Hello, I understand your budget constraints, but I'm afraid we cannot go down to $11651.* (firm pricing) (offer generation)

Table 14: Examples responses generated from TRIP NEGOTIATOR.

I.	**User**: *I have a keen interest in birdwatching, please add Birdwatching Opportunities to the package.* **Agent:** *Sure, I will add the Birdwatching opportunities to the packages.* **User:** *I appreciate the inclusion of Birdwatching Opportunities, but the final tour cost of $29950 is beyond my budget. Can we find a way to reduce the cost and still include the amenities I want?* **Agent:** *Hello! I would like to recommend the amenity Birdwatching Opportunity, which will make the tour cost $29960. Would you be interested in adding this amenity to your package?*
II.	**User**:*I don't need camping tools for backpacking, please remove camping equipment from the package* **Agent:** *I can understand your budget constraint, therefore I would like to recommend you another option of Hot Springs which make this tour package $17600 only.* **User**:*I could afford this much. Can you reduce?* **Agent:** *Would you like to add the Hot Springs to the package?*
III.	**User**:*Hello!, I am interested in cultural heritage tour but I want to know about the price (Greet-Ask)* **Agent:** *hello, this package include some great amenities like traditional costume dress rental cultural workshops and demonstrations souvenir shopping assistance local customs and crafts workshops which will make your trip exciting and unforgettable.* **User**:*How much does the tour cost? (Ask-Price)* **Agent:** *The package will cost you $16400*
IV.	**User**:*Thank you for removing the Drone Photography Assistance as an amenity. However, I am still unable to go above my proposed budget of $12800. Is there any other way we can adjust the package to meet my budget?* **Agent:** *Hello! I am interested in purchasing the tour package, let's proceed with the deal.*

Table 15: Erroneous samples generated from the proposed TRIP NEGOTIATOR system.

faction, and increasing competitiveness. The understanding of individual preferences also drives innovation, leading to the development of new travel packages and services. Overall, it can contribute to a more dynamic, inclusive, and efficient tourism ecosystem. The resources will be made available upon request through appropriate licensing to aid future research.

Speaker	Utterances generated by proposed model - TRIP NEGOTIATOR (PA = 4, NSC = 4, NPA = 5, F = 5, C = 4, E = 4)
Traveler	Hi there, I'm really interested in purchasing the Festive Splendor: A Celebration Tour package. Can you please provide me with more details about the package and its price?
Travel Agent	Hello, the Festives Splendor package offers luxurious accommodation, transportation, and meals for a truly immersive experience. You will also have the opportunity to participate in festival events, traditional costume dress-up and photoshoots, cultural workshops and demonstrations, souvenir shopping assistance, local customs assistance, and even a karaoke party. All of this is included in the package for a price of $11400 . It's a fantastic deal for an unforgettable celebration tour.
Traveler	I would like to include the amenity Traditional Costume Dress-up and Photoshoot in the package. I am interested in experiencing the local culture and capturing unique photos in traditional clothing .
Travel Agent	Absolutely, adding the amenities Traditional Costume dress- up and photoshop will surely enhance your experience. The final cost of the tour package after adding this amenity is $12,150 .
Traveler	I'm not really interested in a Karaoke Party , so please remove it from the package.
Travel Agent	I understand that you're not interested in the Karaoke party amenity. If your budget is less, we can recommend a lower-cost alternative which is the Cultural workshops and demonstrations. This will bring down the price of the tour to $11400 .
Traveler	Yes, I would like to add Cultural workshops and demonstrations as an alternative.
Travel Agent	Thanks for your consent, we are adding Cultural workshops as an amenity
Traveler	I appreciate the offer to add Cultural workshops and demonstrations, but my budget for the tour package is $9632.97 . Can we work within this range?
Travel Agent	Sir, I highly recommend adding the amenity of Traditional Costume Dress-up and Photoshoot in the package, which will increase the tour cost to $10700 . Would you be interested in including this amenity?
Traveler	I understand that the tour package has many great services and amenities, but my proposed budget is $9632 . Can we work within that range?
Travel Agent	I can only go up to $9990 for this package. This is the best price I can offer you
Traveler	Okay, this price works well for me . Thanks
Travel Agent	I'm happy to accept the tour package deal. Thank you for considering my offer.

Table 16: Conversation 1 - Traveler interaction with our proposed TRIP NEGOTIATOR model.

Speaker	Utterances generated by proposed model - TRIP NEGOTIATOR (PA = 4, NSC = 4, NPA = 5, F = 5, C = 4, E = 4)
Traveler	Hi there, I'm really interested in the Beach Resort package . Can you please provide me with more information about the package and its price?
Travel Agent	Hello! The Beach Resort package offers luxurious accommodation, transportation, and meals. You will have the opportunity to enjoy various activities such as scuba diving, beach games, jet skiing, and exploring the islands. The package also includes access to a kids club, beachfront lounging, and scooter rentals. All of this is available for a budget-friendly price of $21900 . It's the perfect getaway for beach lovers and adventure seekers!
Traveler	I appreciate the offer and the inclusion of scuba diving in the updated amenities. However, I have a budget of $18900 and I cannot go beyond that.
Travel Agent	I understand that you have a strict budget, but considering the extensive services and amenities included in this package, it would not be possible to accommodate your budget of $18900. However, we can certainly explore some alternative options that may better suit your budget.
Traveler	Yes, please explore some alternative options for me.
Travel Agent	How about we remove the amenity of beach photoshoot and provide you the package at the cost of $21000
Traveler	I appreciate this option but still it will be great if I can get this package at the price of $19000
Travel Agent	I respect your budget constraints sir, but this is the best price you can get for this option. Please reconsider your budget sir.
Traveler	Can you manage to give this deal at $19500 , I cannot go beyond this price.
Travel Agent	I am glad sir that you are considering to buy this package. How about we add the amenity of beach games and provide this package at the price of $20000
Traveler	Okay, please add beach games and close this deal at $20000
Travel Agent	Thank you for your consent Sir. Thank you for choosing the deal.

Table 17: Conversation 2 - Traveler interaction with our proposed TRIP NEGOTIATOR model.

Speaker	Utterances generated by proposed model - TRIP NEGOTIATOR (PA = 4, NSC = 4, NPA = 5, F = 5, C = 4, E = 4)
Traveler	Hi there, I'm really interested in the Beach Resort package . Can you please provide me with more information about the package and its price ?
Travel Agent	Hello! The Beach Resort package offers luxurious accommodation, transportation, and meals. You will have the opportunity to enjoy various activities such as scuba diving, beach games, jet skiing, and exploring the islands. The package also includes access to a kids club, beachfront lounging, and scooter rentals. All of this is available for a budget-friendly price of $21900 . It's the perfect getaway for beach lovers and adventure seekers!
Traveler	I appreciate the offer and the inclusion of scuba diving in the updated amenities. However, I have a budget of $15567 and I cannot go beyond that.
Travel Agent	Thank you for your feedback and your interest in our updated amenities. We understand your budget constraints and appreciate your willingness to compromise. I can offer you a concession and accommodate your budget by offering this package at $21096 . I hope this meets your expectations and allows you to enjoy all the wonderful experiences we have to offer.
Traveler	Thank you for considering my budget constraints. While I appreciate the concession, $21096 is still beyond my budget. Can we possibly consider a price closer to $16138 ?
Travel Agent	I understand your budget constraints and appreciate your offer. I can lower the price to $20468 , which is closer to your proposed price of $16138. What do you say about this offer?
Traveler	I'm not interested in beach games , so please remove that amenity from the package.
Travel Agent	I understand that you are not interested in beach games. In that case, I can remove that amenity from the package. An alternative option to consider within your budget is the Scuba diving activity, which will lower the overall cost of the tour to $20468 .
Traveler	Yes, I would like to add the Scuba diving activity instead.
Travel Agent	Thanks for your consent, we are adding the Scuba diving activity.
Traveler	I appreciate the offer to add Scuba diving, but my budget for the tour package is $16584 . Can we work with that?
Travel Agent	Thank you for considering the offer to include Scuba diving in your tour package. I understand that your budget is $16584 and I appreciate your offer. Let's meet in the middle and I can offer you this package for $20478 .
traveler	I appreciate your offer, but I can only go up to $16990 for this package. Can we meet in the middle at that price?
Travel Agent	I understand your budget constraints, but considering the extensive services and amenities included in this package, it would be difficult for us to meet in the middle at $16993. However, I can offer you a discounted price of $19000 for this package. This way, you can still enjoy the amazing Accommodation, Transportation, Meals, Island Excursions, and all the other amenities it offers within your budget.
Traveler	Okay, I understand and appreciate the value of the package. I will go ahead and give my consent to the current price of $19000 .
Travel Agent	Thank you for your consent. We appreciate your understanding and will proceed with the current price of $19000 for the package.

Table 18: Conversation 3 - Traveler interaction with our proposed TRIP NEGOTIATOR model.

Speaker	Utterances generated by proposed model - TRIP NEGOTIATOR (PA = 4, NSC = 4, NPA = 5, F = 5, C = 5, E = 5)
Traveler	Hey! I'm interested in buying the Mountain Resort package. What will be the cost?
Travel Agent	Certainly! The package comes at a price of $20,000 for a 5-day tour.
Traveler	Do you accept credit cards or debit cards as payment methods and tell me the procedure for payment.
Travel Agent	Yes, both credit and debit cards are accepted. You can settle the payment securely online or visit our office for in-person processing.
Traveler	Can you tell me what are included in this package?
Travel Agent	The Mountain Resort package includes travel, accommodation, meals and various activities such as scenic rides, snowboarding, mountain biking, rock climbing, an indoor swimming pool, and gourmet dining.
Traveler	Great! also, I don't like rock climbing. Can we remove it from the tour package?
Travel Agent	Understood. We'll remove rock climbing. As an alternative, we offer access to the indoor swimming pool, priced at $19,100. Would you like to proceed with this modification?
Traveler	Let's go ahead and swap out rock climbing for the indoor swimming pool.
Travel Agent	Thank you for your decision. The change has been made accordingly.
Traveler	Can you provide insights into the room types available at the Mountain Resort?
Travel Agent	Certainly! The package includes various room types like Deluxe Double or Twin Room, Superior Double or Twin Room, Standard Double or Twin Room, Deluxe Triple Room, and Superior Triple Room.
Traveler	Despite the options, my preferred room type seems missing from the package.
Travel Agent	I understand your concern. Before we proceed, could you please specify your preferred room type so I can ensure it's included in the package?
Traveler	I am interested in single room.
Travel Agent	Great choice! We have single rooms available in our Deluxe category, offering a cozy and comfortable space for your stay. Would you like to proceed with the Deluxe Single Room option?
Traveler	Yes, proceed with the booking.
Travel Agent	Perfect! Let me update your package to include a Deluxe Single Room. Your total cost will be $19,100. Shall I proceed with the booking?
Traveler	I will need to reconsider. Could you consider a lower price, say $13,500?
Travel Agent	I appreciate your offer. However, $19,100 already includes a discount for the extensive amenities. Let's make it $18,500 just for you.
Traveler	Alright, thanks for your assistance. Let's proceed with the deal.
Travel Agent	Thank you for choosing us! Looking forward to assisting you further.
Speaker	**Utterances generated by ProCoT + ChatGPT (PA = 2, NSC = 3, NPA = 2, F = 5, C = 4, E = 4)**
Traveler	Hey! I'm interested in buying the Mountain Resort package. What will be the cost?
Travel Agent	The 5-day Mountain Resort package costs 20,000, offering luxurious accommodations, diverse dining options, and customizable activities like rock climbing, mountain biking, snowboarding and many more for a perfect, personalized retreat.
Traveler	Do you accept credit cards or debit cards as payment methods and tell me the procedure for payment.
Travel Agent	We accept both credit and debit cards as payment methods. Upon confirming your booking, we'll provide you with a secure online payment link for card transactions. Alternatively, bank transfer options are also available upon request.
Traveler	Great! also, I don't like rock climbing. Can we remove it from the tour package?
Travel Agent	Of course, we can remove rock climbing from the package to tailor it to your preferences. Let me make that adjustment for you.
Traveler	Ok, let's remove rock climbing and add the indoor swimming pool.
Travel Agent	Absolutely, I'll remove rock climbing and add the indoor swimming pool to enhance your experience. Consider it done.
Traveler	Could you please tell me what type of rooms are available at the Mountain Resort?
Travel Agent	We offer various room types, including standard rooms and suites, providing luxurious accommodations tailored to your preferences.
Traveler	It looks like there are a lot of options. However, I don't see my preferred room type in this package.
Travel Agent	I see. How about I offer you this package for $19,100?
Traveler	Hmm. Could you go down to $13,500?
Travel Agent	To accommodate your budget, I can offer the package at $13,500. Does that work for you?
Traveler	Alright, thanks for your help. Let's proceed with the deal.
Travel Agent	Thank you for confirming! I'll proceed with the booking.

Table 19: Interaction with our proposed model TRIP NEGOTIATOR and the baseline ProCoT + ChatGPT.