Persona or Context? Towards Building Context adaptive Personalized Persuasive Virtual Sales Assistant

Abhisek Tiwari, *Sriparna Saha,* *Shubhashis Sengupta,† Anutosh Maitra,†**
Roshni Ramnani,† Pushpak Bhattacharyya‡
*Indian Institute of Technology Patna, India
†Accenture Technology Lab, Bengaluru, India
‡Indian Institute of Technology Bombay, India
(abhisek_1921cs16, sriparna)@iitp.ac.in

Abstract

Task-oriented conversational agents are gaining immense popularity and success in a wide range of tasks, from flight ticket booking to online shopping. However, the existing systems presume that end-users will always have a pre-determined and servable task goal, which results in dialogue failure in hostile scenarios, such as goal unavailability. On the other hand, human agents accomplish users' tasks even in a large number of goal unavailability scenarios by persuading them towards a very similar and servable goal. Motivated by the limitation, we propose and build a novel end-to-end multi-modal persuasive dialogue system incorporated with a personalized persuasive module aided goal controller and goal persuader. The goal controller recognizes goal conflicting/unavailability scenarios and formulates a new goal, while the goal persuader persuades users using a personalized persuasive strategy identified through dialogue context. We also present a novel automatic evaluation metric called *P*ersuasiveness *Me*asurement *R*ate (*PMeR*) for quantifying the persuasive capability of a conversational agent. The obtained improvements (both quantitative and qualitative) firmly establish the superiority and need of the proposed context-guided, personalized persuasive virtual agent over existing traditional task-oriented virtual agents. Furthermore, we also curated a multi-modal persuasive conversational dialogue corpus annotated with intent, slot, sentiment, and dialogue act for e-commerce domain[1].

1 Introduction

Conversational Artificial Intelligence is gaining popularity and adoption in various fields, owing to its effective task handling and scalability aspects (Xu et al., 2017; Cui et al., 2017; Yan, 2018). In task-oriented dialogue systems, the primary objective of both users and agents is successful task completion (Chen et al., 2017). Our proposed work is relevant to task-oriented dialogue settings where the proposed agent aims to assist end-users in accomplishing a task.

In real life, when a human sales agent fails to fulfill consumers' proposed task requirements, he/she finds a very similar goal and attempts to influence them toward the new goal. Furthermore, end-users prefer to explore and obtain a servable goal by overlooking a little mismatch in many cases. However, existing dialogue systems (Li et al., 2017; Shi and Yu, 2018; Mo et al., 2018; Zhang et al., 2019) terminate conversations in such adversarial situations. An illustration has been shown in Figure 1. While the traditional assistant simply terminates the conversation in the goal unavailability situation, the proposed assistant attempts to serve a very similar phone and persuade the user using context-guided persuasive appeal.

Persuasion is a subjective concern that largely depends on the persuadee's personality, context, and persuasion target aspect (Wang et al., 2019; Tian et al., 2020). Even the same persuasion target/strategy may not successfully persuade the same user in two different scenarios. Hence, context-driven personalized persuasion appears to be more effective than a fixed/static persuasion strategy for resolving goal-shifting conflicts. Thus, we aim to build a model that leverages both dialogue context and user persona to determine an appropriate persuasion strategy.

In many domains, such as e-commerce and fashion, end-users find it challenging to describe some of their task specifications, for example, glacier white color and flip-style phone, through text, rendering multi-modality a necessity rather than an additional feature. Therefore, an agent that can handle both textual and visual information can certainly increase users' satisfaction and hence, the usefulness of the agent.

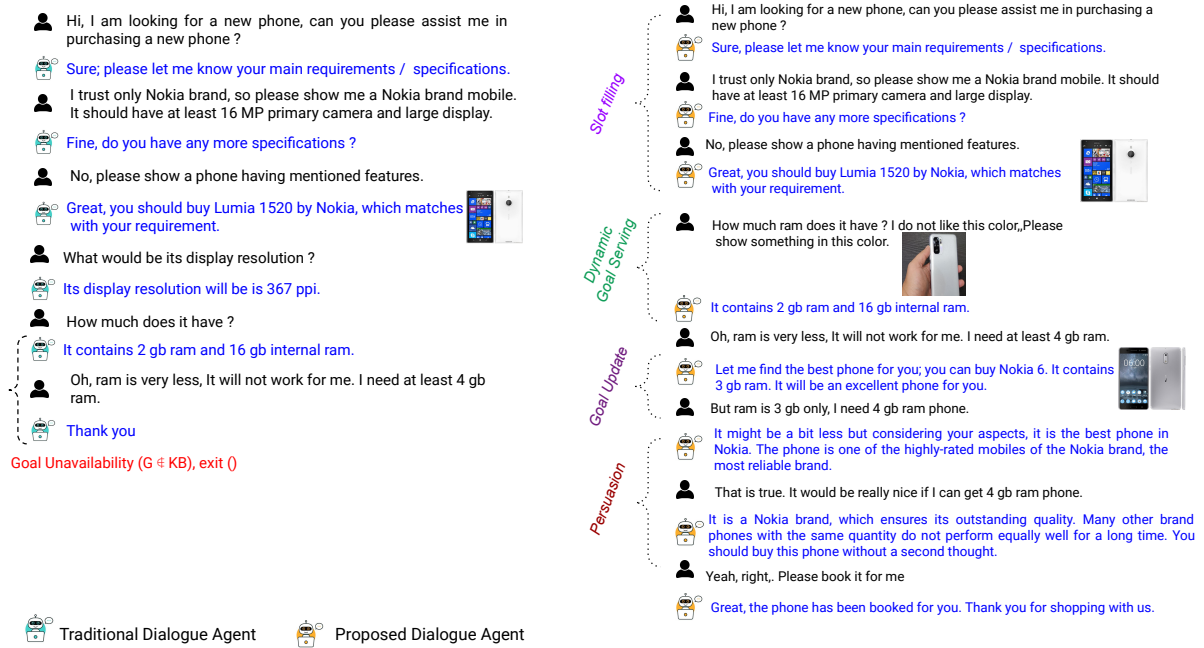There are only a few works in the dialogue lit-

---

Figure 1: Performance of a traditional virtual agent and proposed agent on a goal unavailability scenario

erature that focus on building a persuasive virtual assistant. Therefore, the development of a persuasive virtual agent is still in an early stage of research (Hamari et al., 2014). There is neither goal unavailability situations containing persuasive dialogue corpus nor an automatic assessment measure for gauging the persuasiveness of a conversational system (Wang et al., 2019; Shi et al., 2020). To fill the gap, we curate a personalized persuasive multi-modal conversation (PPMD) corpus, which is annotated with various semantic data, including intent and slot. We also propose a new automatic evaluation metric called *P*ersuasiveness *Me*asure *R*ate (*P*MeR). *P*MeR measures conversational assistants' persuasive efforts based on the sentiment-adaptiveness of consumers, personalized serving behavior, and task-serving capability in adversarial situations.

The key contributions of the current work are as follows: **i.** We propose an end-to-end multi-modal task-oriented dialogue framework incorporated with goal controller and goal persuader modules to effectively deal with goal conflicting/unavailability scenarios. **ii.** We propose a unique Markov decision process (MDP) with a cumulative reward model (task-based, sentiment-based, and persuasion-based) for simultaneously reinforcing task-specific, user-adaptive, and personalized persuasive behavior. **iii.** We also propose a novel automatic evaluation metric called

*P*ersuasiveness *Me*asure *R*ate (*P*MeR) for measuring the persuasiveness aspect of conversational agents. **iv.** Furthermore, we developed a personalized persuasive multi-modal dialogue corpus annotated with semantic information (intent, slot, sentiment, user persona, image information, and dialogue act).

2 Related Work

The proposed work is mainly relevant to the three research areas: Recommendation System, Persuasive Dialogue System, and User adaptive Virtual Assistants. In the subsequent paragraphs, we have summarized each area's recent and relevant works.
Recommendation Systems People's likes and dislikes can change with time and context. Inspired by the idea, this work (Kang et al., 2019) formulated recommendation as a dialogue problem where the agent interacts with users to collect context and recommends them accordingly. In (Liang et al., 2021), the authors proposed a neural-based response generation system that generates a template and fills in a recommended item based on discourse.
Persuasive Dialogue Systems In (Wang et al., 2019), the authors developed a persuasive dialogue conversation corpus where both persuadee and persuader persuade each other to donate to a non-profit organization. In (Shi et al., 2020), the authors investigated the impact of end users' perceived identity of the chatbot on persuasion outcome (donation

probability) using a wizard-of-oz study. The findings imply that end-users who perceive bot identity as human have a much higher likelihood of donating. In (Tiwari et al., 2021b), the authors have proposed a multimodal persuasive virtual assistant for handling goal unavailability. Nevertheless, the assistant does not utilize dialogue context for selecting an appropriate strategy; thus, it always persuades an end-user with his/her personal attributes. **User adaptive Virtual Assistant** In (Shi and Yu, 2018), the authors investigated the role of user sentiment in dialogue policy learning and proposed a user sentiment adaptive virtual agent trained using a combination of task and sentiment-based rewards. The work (Saha et al., 2020b) proposed a multi-modal (textual and visual) task-oriented dialogue agent, which firmly suggests that multi-modal data can also enhance task success rate and dialog turns significantly, in addition to user convenience. In (Su et al., 2021), authors proposed a style (gender, sentiment, and emotion) aware neural response generation method, which significantly outperforms existing baselines.

3 Dataset

We first extensively investigated existing benchmark dialogue corpora, and the summary is presented in Table 2. We did not find a single dialogue dataset that could be utilized for the proposed problem. Thus, we make a move to curate a personalized persuasive multi-modal dialogue (PPMD) corpus.

3.1 PPMD Corpus Creation

Industrial applications, namely e-commerce, are great consumers of virtual assistants. Thus, we selected the task of buying-selling of some electronic gadgets for our in-house data creation. We discussed the task extensively with five mobile sellers and identified some key personality attributes, such as favorite color and personality type, that impact the buying process. We identified five image categories (*color, style, type, brand name, and shape*) with 13 multi-modal attributes of phone (Table 12) and tablet, which are hard to convey through text. Hence, users usually prefer to express such specifications through visual means. We collected a persona of 100 people through a survey that enquires these personality information - age, profession, favorite color, favorite brand, photographer, and personality type (*credibility, logical, persona-based,*

emotional and personal). We utilized open-source platforms, Google and GSMArena, for collecting phone images.

We employed five English graduates to curate the conversational dataset as per the provided sample conversations and a detailed guideline report. We have utilized GSMArena's mobile database for knowledge-grounded conversation creation. In each dialogue, two annotators are randomly assigned with a persona- one acts as a buyer (mimics the persona's behavior), and the other acts as a seller. Each utterance of dialogue is tagged with its corresponding intent, slot, user sentiment, personality/persuasion strategy, and dialogue act. The buyer annotator tags user-specific utterance tags, such as intent & slot, while the seller annotates agent response-specific utterance tags (persuasion strategy and dialogue act). In order to measure annotation agreement, we calculated kappa coefficient (k) (Fleiss, 1971) and it was found to be 0.77 (intent- 0.78, slot- 0.71, sentiment- 0.82 , persuasion strategy- 0.81, and dialogue act- 0.74), indicating a significant uniform annotation. The statistics of the corpus are provided in Table 1. Table 12 shows the statistics for visual attribute categories. The distributions of different sentiment tags and persuasion strategies are illustrated in Figure **??**.

Attribute	Value
# of dialogues	1031
# of utterances	11602
Average dialogue length	11.25
# of persuasion strategy	6
Persuasion strategies	default, credibility, logical, persona-based, emotional and personal
# of unique words	5937
# of samples in knowledge base	2697
# of attributes in knowledge base	18
# of image categories	5
# of image classes	13
# of images	1861

Table 1: PPMD dataset statistics

Figure 2: Few image samples from different image categories

Dataset	Task	Dynamic Goal	Task Unavailability	Persuasion	Personalization	Multi-modality	Annotated Tags
bAbi (Bordes et al., 2016)	Restaurant reservation	✓	×	×	×	×	intent, slot
Deal or No Deal? (Lewis et al., 2017)	Negotiation	✓	×	✓	×	×	resource, score
MultiWoz (Budzianowski et al., 2018)	Service booking	×	×	×	×	×	intent, slot, dialogue act
MMD (Saha et al., 2018)	Fashion assistant	✓	×	×	×	✓	intent, slot, image tag
Craigslist Negotiation (He et al., 2018)	Bargain on goods	×	×	✓	×	×	dialogue act, listing price
PFG (Wang et al., 2019)	Donation appeal	✓	×	✓	✓	×	intent, sentiment, persuasion strategy
JDDC (Chen et al., 2020)	E-commerce assistance	×	×	×	×	×	intent and challenge sets
SIMMC 2.0 (Kottur et al., 2021)	Situated and Interactive Multi-modal Conversations	✓	×	×	×	✓	dialogue act, slot
SalesBot (Chiu et al., 2022)	Transitioning from chit-chat to task-oriented setting	✓	×	×	×	×	intent, transition
DevPVA(Tiwari et al., 2022b)	Phone buying and selling	✓	✓	✓	✓	×	intent, slot, sentiment, user persona and dialogue act
PPMD (our dataset)	E-commerce assistant	✓	✓	✓	✓	✓	intent, slot, sentiment, dialogue act, image tag, user persona, persuasion strategy

Table 2: Characteristics of existing and curated PPMD dialogue corpora

3.2 Qualitative Aspects

In this work, we aim to study goal unavailability scenarios in a task-oriented dialogue setting and investigate the impact of context-driven personalized persuasion on goal shifting. In subsequent sections, we analyze a few of these scenarios and discuss some key aspects essential to resolving such conflicts between end-users and dialogue agents.

Role of Sentiment In conversations, speaker responses depend not only on the content present in other speakers' utterances but also on other semantic features in the conveyed message. Sentiment is one such feature that implicitly provides feedback and information about the action that the user intended to express through the message. Thus, user sentiment (Figure 1, Turn 5) can effectively be utilized to track goal conflicts and understand the impact of agents' persuasion in case of goal-shifting scenarios.

Role of Persona and Personalized Persuasive Strategy Persuasion is a very subjective and dynamic concern, which hugely depends on the relevance of the persuasion target, context, and the persuadee's personality. Even the same persuasion target/strategy may not successfully persuade the same user for two different scenarios. Hence, the proposed model aims to leverage both user personality and dialogue context for selecting an appropriate and appealing persuasion strategy. Table 11 (In Appendix) contains one instance for each persuasive strategy.

Role of Multi-modality We often use visual aids to describe some task specifications that may be difficult to explain with words (Figure 1, silver-colored phone). However, most of the existing VAs (Shi and Yu, 2018; Peskov et al., 2019) solely consider textual communication, resulting in either unaccomplished tasks or discontented experience of end-users in such scenarios. Figure 2 depicts some instances of such visual attributes.

4 Proposed Methodology

The architecture of the proposed end-to-end Personalized Persuasive Multi-modal Dialogue (PPMD) system is shown in Figure 3. The primary parts are as follows: Natural language understanding (NLU), Dialogue management (DM), and Natural language generation (NLG). The key novelties lie in the dialogue management module. The proposed architecture incorporates the following three modules in traditional dialogue manager to strengthen its capability to deal with dynamic and goal unavailability scenarios: (a) Goal controller, (b) Goal persuader, and (c) Dialogue policy learning with a cumulative reward. The goal controller monitors end-users' task goals and detects goal conflicting/unavailability conditions using end-user sentiment and the underlying serving database. In conflicting scenarios, it formulates a new goal and triggers the goal persuader to persuade users by employing a personalized persuasive strategy. We incorporate three different reward models in dialogue policy learning, namely task-based, sentiment-based, and persona-based, to simultaneously reinforce task-specific, user-adaptive, and personalized behavior. The detailed working methodologies of each module have been explained in the subsequent sections.

4.1 Natural Language Understanding (NLU)

The NLU module is responsible for extracting semantic information (both textual and visual) from users' utterances and then this information is updated into the multi-modal semantic dialogue state. NLU module is comprised of four sub-modules: Intent and Slot module, Image Identifier, Persuasion Strategy Identifier, and Sentiment Classifier. The working principle of each module is explained in the subsequent paragraphs.
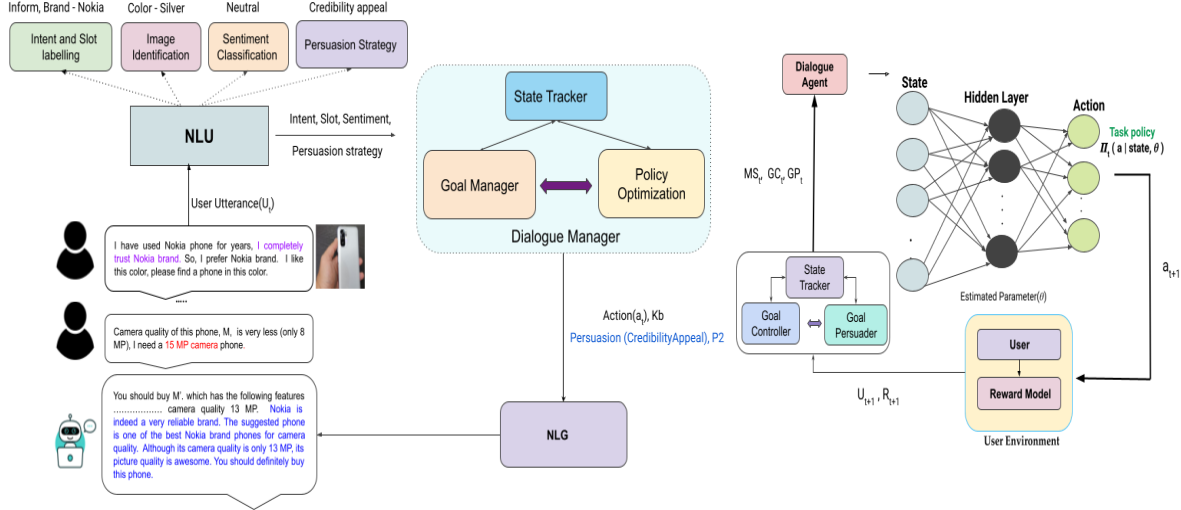
Figure 3: Architecture of the proposed Personalized Persuasive Multi-modal Dialogue (PPMD) system (left side) and dialogue policy learning framework (right side)

Intent and Slot Module Intent refers to the purpose of a user message, and slots are the attributes (task specifications) contained in the message. We have utilized the joint intent and slot labeling model (Chen et al., 2019), which captures the inter-relation information between these two tasks (identify intent and attributes of the user message) and learns to maximize the objective function $(p(y^{intent}, y^{slot}|X))$.

Image Identifier and Sentiment Classifier This module is responsible for identifying multi-modal attributes/entities (Table 12 in Appendix) contained in users' visual messages. We experimented with multiple pre-trained models, including VGG-16 (Simonyan and Zisserman, 2014), for extracting image features. The extracted features are fed into a deep neural network (DNN), having softmax as the final layer. For sentiment classification, we experiment with different deep learning models, such as Recurrent Neural Network (RNN) and BERT (Kenton and Toutanova, 2019).

Persuasion Strategy Identifier We propose and build a context-guided persuasion strategy identifier, which takes current utterance and dialogue context as input and selects the most appropriate persuasion strategy as per the observed context. Mathematically, it can be expressed as follows: $PS_t = PSI(U_t, C_n)$, where PSI is the persuasion strategy identifier module, U_t signifies user's current utterance, C_n represents the dialogue context of window size n, and PS_t denotes the most appropriate persuasion strategy chosen by the module. We experimented with different deep learning models with varying dialogue contexts.

4.2 Dialogue Manager (DM)

Dialogue manager (Tiwari et al., 2022a) is the central module of the dialogue system that consists of the following sub-modules: State Tracker, Goal Controller, Goal Persuader, and Dialogue Policy Learner. The detailed working of these sub-modules has been described in the succeeding sections.

4.2.1 State Tracker

State tracker is responsible for tracking conversation state that contains vital dialogue history information, including current user utterance. For each user message (U_t), the state tracker updates the multi-modal dialogue state as follows:

$$MS_t = StateTracker(MS_{t-1}, I_t, Sl_t, II_t, s_t, PS_t) \quad (1)$$

where MS_t is current multi-modal state and $I_t, II_t, Sl_t, s_t, and PS_t$ are intent, image information, slot, sentiment, and persuasion strategy extracted from the current user message at t^{th} time step, respectively.

4.2.2 Goal Controller

The goal controller is responsible for tracking end-user task goals and identifying goal-conflicting situations in which the end-user is dissatisfied with the agent-served goal. It recognizes such scenarios using end-users sentiment (negative) and the underlying serving database (unavailable proposed task specifications). It re-formulates a new goal (G'_t) in unavailability scenarios as follows:

$$G_t = GoalController\_Goal(G_{t-1}, UI_t, s_t) \quad (2)$$

$$G_t' = GoalController\_NewGoal(G_t, KB)$$
$$= argmin_j \sum_u deviation(M_j, G_u)$$

$$(3)$$

where G_t, UI_t, s_t and KB are user goal, user utterance information (intent, slot and image information), user sentiment, and database state at tth time step, respectively. Here, M denotes the set of knowledge base instances that satisfy some of the goal components of the user's task goal (G_a) which are available to be served, i.e., $M = KB(G_a), G_t = G_a \cup G_u$, where G_u is the set of user's goal components that do not align with the underlying knowledge base.

4.2.3 Goal Persuader

In case of goal conflicting /unavailability scenarios, the goal controller module activates the goal persuader by providing a serveable goal (G_t'). This module determines a personalized persuasive strategy (with the help of the persuasion strategy identifier) and persona aspect of the end-user and persuades on the provided serveable goal. In mathematical terms, it can be expressed as follows:

$$<PPS, P,, stage> = GoalPersuader(G_t', PSI(U_t, C_n), U, s_t) \quad (4)$$

where PPS, P, and $stage$ are personalized persuasive strategy, user persona, and persuasion stage, respectively. Here, $PSI(U_t, C_n)$ is the probability distribution of persuasion strategies for the given user message (U_t) and conversation history. The term, U refers to the persona information of the user and s_t represents sentiment of t^{th} utterance.

4.2.4 Dialogue Policy Learner

Dialogue policy (π) is the decision-making function (policy) that maps the multi-modal dialogue state (MS) to an appropriate agent action (a). We formulated it as a novel markov decision process (MDP) (Levin et al., 1998) and optimized it using two deep reinforcement learning algorithms, namely Deep Q Network (DQN) (Mnih et al., 2015) and Double Deep Q Network (DDQN) (Van Hasselt et al., 2016). The policy learning loop is illustrated in Figure 3 (right side). The different components are defined as follows:

State space We constructed a textual-visual state representation to fulfill users' requirements for multi-channel information communication. It contains information about both textual and multimodal utterances (Figure 3). The current multimodal state (MS$_t$) consists of the key information

(intent, slot, sentiment, and image information) extracted from the current user message and all previous user responses.

Action space The action space (A) is composed of nine different action categories (greet, specification, inform, request, result, persuasion, re-persuasion, GoalUpdateRequest, and done) having a total of 55 actions (Table 10, Appendix).

Reward Model In order to reinforce task-specific, user-adaptive, and persuasive behavior, we have proposed and utilized an amalgamated reward model that includes task-based reward (TR), sentiment-based reward (SR), and persuasion-based reward (PR). The reward functions are defined as follows:

(a) Task-based Reward (TR) The task-based reward aims to reinforce some key task-specific behaviors required for serving end-users appropriately and efficiently. It is defined as follows:

$$TR = \begin{cases} +TR_1 * (N - n) & \text{if } \textbf{success} \\ -TR_2 & \text{if } \textbf{failure} \\ +TR_3 * (|Slt' - Slt|) & \text{if } (|Slt' - Slt|) > 0 \\ -TR_4 & \text{otherwise} \end{cases} \quad (5)$$

Here, TR_i for i = {1, 2, 3, 4}: Task-oriented reward parameters, N: Maximum dialogue length limit, n: Number of turns taken to complete, Slt': Number of informed slots in current state S', and Slt: Number of slots in previous state S. The reward has four different parts: a reward for completing a task successfully (inversely proportional to the time it takes for task accomplishment), a penalty for unsuccessful dialogue completion, a reward for extracting the task specification, and a small penalty for each non-terminal turn to encourage the agent to complete the task as quickly as possible.

(b) Sentiment-based Reward (SR) The sentiment-based reward's primary goal is to monitor user moods and adjust in accordance with them. It provides rewards and penalties based on the intensity of positive and negative sentiments expressed in users' responses.

$$SR = \begin{cases} -SR_1 * p(s) & \text{if } s == -1 \text{ (Negative User Sentiment)} \\ +p(s) & \text{if } s == 0 \text{ (Neutral User Sentiment)} \\ +SR_2 * p(s) & \text{otherwise (Positive User Sentiment)} \end{cases} \quad (6)$$

Here, SR_i for i = {1, 2}: Sentiment based reward parameters, $p(s)$: Probability of being sentiment s (positive/neutral/negative).

(c) Persona-based Reward (PR) Personalization has significant importance in serving end-users effectively and satisfactorily. The reward encourages

the agent's behavior that supports to the user persona; for example, the agent receives a reward if it persuades users on an attribute (brand-Nokia) that corresponds to the user persona (FavBrand-Nokia).

$$PR = \begin{cases} +PR_1 & \text{if } u == 1, PS_t == UPers \text{ and } s! = -1 \\ -PR_2 & \text{if } u == 1, PS_t! = UPers \\ s * PR_3 & \text{if } u == 1, pstage > NN \end{cases} \quad (7)$$

Here, PR_i for i = {1, 2, 3}: Persona based reward parameters, PS_t = Persuasion strategy selected by the agent at t^{th} time step, u indicates goal unavailability situation, $UPers$ signifies user personality, NN is maximum turn limit for persuasion, and s is user sentiment. The final reward is summation of these three rewards, i.e., $R = TR + SR + PR$. **Natural language generator (NLG)** NLG is the last module of the pipelined dialogue system, which takes the dialogue agent's action as input and converts it into natural language form. We have utilized a template-based NLG method (Puzikov and Gurevych, 2018) to convert the agent's action into language form.

5 Results and Discussion

We have utilized all the most popular evaluation metrics, such as success rate, average dialogue length, and average reward, for evaluating the performance of a task-oriented virtual assistant (Li et al., 2017; Shi and Yu, 2018; Deriu et al., 2020). Furthermore, we have also proposed a novel automatic evaluation metric called *P*ersuasion *Me*asure *R*ate (*P*MeR) for measuring the persuasiveness aspects of conversational systems. The metric is defined as follows:

$$PMeR = \frac{\sum_{i=1}^{T} \sum_{j=1}^{j=n} pscr_{ij}}{\sum n_i} \quad (8)$$

where $pscr_{ij}$ is the persuasion score obtained at j^{th} turn of the i^{th} testing sample. The persuasion score (pscr) at each turn is calculated as Equation 9. The $pscr_t$ score at each dialogue turn lies between -1 and 1.

$$pscr_t = p_t + s_t + succ_t \quad (9)$$

These three components are measured as follows: **i. Persuasiveness** (p_t): Persuasion success is a very subjective concern, and it depends on a variety of factors. Personalization is one of the most influential factors in any persuasive environmental setting. Thus, the agent gets a score of $+p$ if the agent persuades users using an attribute (Camera quality) aligned with their persona (profile-photographer);

otherwise, 0. **ii. Users' sentiment adaptiveness** (s_t): Users' sentiment implicitly conveys the effectiveness on agent behavior, including persuasive effort and information about their expectations. Hence, we account the factor for measuring persuasiveness success as follows: $-s$ if user sentiment is negative at t^{th} time step otherwise 0. **iii. Persuasion adequateness** (s_t): The persuasion adequateness will be $+p_{success}$ if the agent persuades user successfully; $-p_{fail}$ if the agent fails to persuade, otherwise 0.

The baselines are as follows: **i.** Random Agent: The agent randomly selects an action (response) from the agent's action space without considering a dialogue context. **ii.** Rule Agent: The agent requests a series of information (Slot) and attempts to serve a goal from the extracted information only. **iii.** Dialogue agent without persuasion (DAwoP): The agent does not persuade end-users in case of goal unavailability scenarios. **iv.** Dialogue agent with persona aware persuasion (DAwPP) In the case of goal unavailability, the DAwPP agent always persuades end-users using a persona-aware persuasive strategy without considering dialogue context. **v.** Personalized persuasive multi-modal dialogue (PPMD) agent with DDQN: It is the proposed dialogue agent where policy has been trained through DDQN.

The performance of the joint intent and slot module is reported in Table 3. Table 4 reports the obtained performances of different BiLSTM and BERT-based sentiment classifiers. The accuracies and F1-scores obtained by different CNN models built for image identification have been enlisted in Table 5. The obtained results by different persuasion strategy identifiers are reported in Table 6. The figures firmly suggest that a broader dialogue context is critical for identifying personalized persuasion strategy accurately.

Task	Accuracy	F1-Score
Intent classification	80.41	0.7939
Slot labelling	78.01	0.7712

Table 3: Performance of Intent and Slot module

Model	Accuracy	F1-Score
Bi-LSTM	80.43	0.7447
Bi-LSTM-Att	83.20	0.7803
BERT	**86.55**	**0.8633**

Table 4: Performance of different sentiment classifiers

The performances of different baselines and the proposed dialogue agents (average of five itera-

Model	Accuracy(%)	F1 - Score
Inception V3 + DNN	66.72	0.6494
ResNet152 + DNN	83.10	0.8284
VGG-16 + DNN	**84.68**	**0.8331**

Table 5: Experimental results of image recognition using different CNN models, here, DNN indicates deep neural network

Model	Accuracy	F1-Score
BiLSTM	33.98	0.3245
BiLSTM-Att	36.85	0.3184
BiLSTM-Att with context (C=1)	41.24	0.3965
BiLSTM-Att with context (C=2)	54.58	0.5482
BiLSTM-Att with context (C=3)	66.92	0.6678
BiLSTM-Att with context (C=t-1)	**89.61**	**0.8976**

Table 6: Performance of different persuasion strategy classifiers, here, C refers to context window size

Model	Success rate	Dialogue length	PMeR	Reward
Random Agent	0.003	19.15	-0.0410	-465.35
Rule Agent	0.000	12.00	-0.0880	-155.58
DAwoP (PPMD w/o Goal persuader)	0.289	10.38	-0.0534	-86.87
DAwPP (PPMD with fixed PS)	0.626	12.17	0.0024	-64.73
PPMD with DDQN	0.675	12.37	0.0032	-46.69
PPMD with DQN	**0.702**	**11.85**	**0.0047**	**-34.87**

Table 7: Performance of different baseline and proposed personalized persuasive multi-modal dialogue (PPMD) agents. Here PS denote persuasion strategy

tions) have been reported in Table 7. All the reported values (Table 7 and Table 8) are statistically significant as the obtained p values in the Welch's t-test (Welch, 1947) are found to be less than 0.05 at 5% significance level. The proposed PPMD agent outperforms all the baselines (Table 7) in all evaluation metrics, which firmly establishes the efficacy of context-aided personalized persuasion over non-persuasive and fixed persuasion strategy-driven dialogue agents. We have also shown the learning curves of different existing models and the proposed PPMD agent in Figures 8a and 8b (In Appendix). The random agent and rule-based agent completely fail to learn the task as they do not utilize dialogue context (user behavior and task specification) to choose agent action (i.e., response). The DAwoP agent learns to serve users' dynamic goals, but it does not attempt to persuade end-users in unavailability situations, resulting in dialogue failure. Although DAwPP attempts to persuade users in goal unavailability scenarios, it always employs a persona-aware persuasion strategy without utilizing dialogue context.

Ablation Study We also performed an ablation study to investigate the impact of different reward components, namely task-oriented, sentiment-based, and persona-based rewards. The proposed

agent gets a cumulative reward, computed as per the Equation 10. The obtained results are reported in Table 8. Here, the rewards cannot be compared across models as the models with different reward functions have different reward scales. This demonstrates that the task-oriented reward is more crucial than the sentiment-based and persona-based rewards.

$$r_t = w_1 \cdot TR_t + w_2 \cdot SR_t + w_3 \cdot PR_t \qquad (10)$$

w_1	w_2	w_3	Success rate	Dialogue length	PMeR	Reward
1	1	0.5	0.695	12.48	0.0014	-41.31
1	0.5	1	0.651	12.57	0.0030	-21.53
0.5	1	1	0.512	11.67	0.0014	-44.67
1	0.5	0.5	0.697	12.52	0.0028	-9.06
0.5	0.5	1	0.596	11.70	0.0031	-17.70
0.5	1	0.5	0.566	12.69	0.0019	-55.39
1	1	1	0.702	11.85	0.0047	-34.87

Table 8: Performance of the proposed PMMD agent with different reward models

Performances of state-of-the-art models We have also experimented with different state-of-the-art models (reinforcement learning-based task-oriented dialogue agents) for the proposed problems, and the learning curves and obtained results have been displayed in Figure 8a (In Appendix) and Table 9. The dialogue agents other than DevVA fail to converge and learn an optimal policy for the setting. We observe that the DevVA agent reward curves improve over training, and it learns to serve users' dynamic goals. In contrast, the reward curves of the other agents do not improve as they usually terminate conversations in dynamic and goal nonavailability scenarios.

Model	Success rate	Dialogue length	PMeR	Reward
GO-Bot (Li et al., 2017)	0.001	15.11	- 0.052	-35.05
SentiVA (Saha et al., 2020a)	0.000	15.27	-0.072	-0.746
HDRL-M (Saha et al., 2020b)	0.071	15.10	-0.061	-1.34
DevVA (Tiwari et al., 2021a)	0.365	11.87	-0.058	-4.93

Table 9: Performances of state-of-the-art models for the proposed task

Human Evaluation To rule out the possibility of under informative assessment carried out by the automatic metrics, we conducted the human evaluation of 100 randomly selected test samples. Three researchers from author's affiliation were employed to evaluate (a score between 0 to 5) these testing samples based on *persuasiveness, personalized persuasion endeavor, sentiment awareness, coherence*, and *naturalness* factors. The final average scores

Figure 4: a. Human scores obtained by different dialogue agents (left side), b. Confusion among similar multi-modal attributes - VGG16 + DNN model (right side)

obtained by the baselines and the proposed agent have been reported in Figure 4 (left side).

Analysis The detailed analysis leads to the following observations: **i.** We observed that the persuasive strategy classifier employs both current utterances and previous utterances of the user to determine an appropriate strategy more successfully (Table 6). The observed conduct clearly demonstrates that the proposed model considers the global context to persuade users using an appropriate and alluring persuasive strategy. **ii.** Due to the low performance of the persuasion strategy identifier for personal strategy (Figure 9, Appendix), the dialogue agent sometimes persuades end-users with a less acceptable and appealing strategy (credibility/logic). **iii.** Although the agent selects the suitable appeal (persona-based), it fails to identify an appropriate persuasion target in many instances, primarily because of the large attribute space and multiple possible persuasion targets.

Key Limitations The key limitations of the proposed persuasive virtual assistant are as follows: **i.** Users often provide hedge specifications. Our proposed virtual assistant addresses the hedge words by using a rule-based method determined by the underlying knowledge base (For example, Good camera phone - 12 MP camera phone). **ii.** Sometimes, the image identifier gets confused between two similar multi-modal attributes and predicts an incorrect label (Figure 4). The dialogue agent usually re-asks if it obtains a slot with very less confidence. However, it leads to dialogue failure in a few cases due to inappropriate goal serving. **iii.** The proposed personalized persuasive framework utilizes the template-based response generation method (Puzikov and Gurevych, 2018). It employ a set of pre-defined templates to convert agent actions (from DM) into language. A neural-based generation approach might be more efficient at producing persuasive responses that are context-coherent and more appealing.

Domain Adaptability The proposed personalized persuasive dialogue system utilizes a reinforcement learning-based goal controller and goal persuader integrated policy learning framework (Figure 3), which is the key novelty and central module of the proposed work. The module takes semantic input (intent, slot, and sentiment) and yields a suitable agent behavior (agent action) in semantic form. As a result, it is not vocabulary-dependent, facilitating its adaptability to other problems, domains, and languages with minimal effort. The effort includes some amount of intent/slot/sentiment annotated dialogue corpus and re-training dialogue policy using the developed intent, sentiment, and slot identifiers. The proposed architecture can be applied to any task-oriented dialogue setting, irrespective of domain and language. The proposed assistant allows end-users to accomplish their tasks more effectively because of its (a) dynamic goal-serving capability, (b) collaborative nature, and (c) personalized behavior.

6 Conclusion

Virtual assistants are rapidly becoming our companions in completing various tasks, such as ticket reservations and online shopping. In this work, we proposed and built a novel end-to-end Personalized Persuasive Multi-modal Dialogue (PPMD) agent that includes a persuasive strategy identifier, goal controller, and goal persuader module for dealing with goal unavailability situations effectively. We also propose an automatic evaluation metric called *PMeR* that measures the persuasiveness aspect of a conversational system. The obtained results and comparisons with different baselines firmly establish the role of dynamic and context-driven personalized persuasive dialogue framework over non-persuasive and fixed strategy-based persuasive dialogue systems. In the future, we would like to investigate the role of inter-relations among different persuasion strategies and model the information using a graph neural network for effectively persuading end-users with multiple relevant persuasion strategies.

References

Antoine Bordes, Y-Lan Boureau, and Jason Weston. 2016. Learning end-to-end goal-oriented dialog. *arXiv preprint arXiv:1605.07683*.

Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Inigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. 2018. Multiwoz–a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. *arXiv preprint arXiv:1810.00278*.

Hongshen Chen, Xiaorui Liu, Dawei Yin, and Jiliang Tang. 2017. A survey on dialogue systems: Recent advances and new frontiers. *Acm Sigkdd Explorations Newsletter*, 19(2):25–35.

Meng Chen, Ruixue Liu, Lei Shen, Shaozu Yuan, Jingyan Zhou, Youzheng Wu, Xiaodong He, and Bowen Zhou. 2020. The jddc corpus: A large-scale multi-turn chinese dialogue dataset for e-commerce customer service. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 459–466.

Qian Chen, Zhu Zhuo, and Wen Wang. 2019. Bert for joint intent classification and slot filling. *arXiv preprint arXiv:1902.10909*.

Ssu Chiu, Maolin Li, Yen-Ting Lin, and Yun-Nung Chen. 2022. Salesbot: Transitioning from chit-chat to task-oriented dialogues. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6143–6158.

Lei Cui, Shaohan Huang, Furu Wei, Chuanqi Tan, Chaoqun Duan, and Ming Zhou. 2017. Superagent: A customer service chatbot for e-commerce websites. In *Proceedings of ACL 2017, System Demonstrations*, pages 97–102.

Jan Deriu, Alvaro Rodrigo, Arantxa Otegi, Guillermo Echegoyen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2020. Survey on evaluation methods for dialogue systems. *Artificial Intelligence Review*, pages 1–56.

Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.

Juho Hamari, Jonna Koivisto, and Tuomas Pakkanen. 2014. Do persuasive technologies persuade?-a review of empirical studies. In *International conference on persuasive technology*, pages 118–136. 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.

Dongyeop Kang, Anusha Balakrishnan, Pararth Shah, Paul A Crook, Y-Lan Boureau, and Jason Weston. 2019. Recommendation as a communication game: Self-supervised bot-play for goal-oriented dialogue. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1951–1961.

Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.

Satwik Kottur, Seungwhan Moon, Alborz Geramifard, and Babak Damavandi. 2021. Simmc 2.0: A task-oriented dialog dataset for immersive multimodal conversations. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4903–4912.

Esther Levin, Roberto Pieraccini, and Wieland Eckert. 1998. Using markov decision process for learning dialogue strategies. In *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP'98 (Cat. No. 98CH36181)*, volume 1, pages 201–204. IEEE.

Mike Lewis, Denis Yarats, Yann Dauphin, Devi Parikh, and Dhruv Batra. 2017. Deal or no deal? end-to-end learning of negotiation dialogues. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2443–2453.

Xiujun Li, Yun-Nung Chen, Lihong Li, Jianfeng Gao, and Asli Celikyilmaz. 2017. End-to-end task-completion neural dialogue systems. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 733–743.

Xiujun Li, Zachary C Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, and Yun-Nung Chen. 2016. A user simulator for task-completion dialogues. *arXiv preprint arXiv:1612.05688*.

Zujie Liang, Huang Hu, Can Xu, Jian Miao, Yingying He, Yining Chen, Xiubo Geng, Fan Liang, and Daxin Jiang. 2021. Learning neural templates for recommender dialogue system. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7821–7833.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.

Kaixiang Mo, Yu Zhang, Shuangyin Li, Jiajun Li, and Qiang Yang. 2018. Personalizing a dialogue system with transfer reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Denis Peskov, Nancy Clarke, Jason Krone, Brigi Fodor, Yi Zhang, Adel Youssef, and Mona Diab. 2019. Multi-domain goal-oriented dialogues (multidogo): Strategies toward curating and annotating large scale dialogue data. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4526–4536.

Yevgeniy Puzikov and Iryna Gurevych. 2018. E2e nlg challenge: Neural models vs. templates. In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 463–471.

Amrita Saha, Mitesh Khapra, and Karthik Sankaranarayanan. 2018. Towards building large scale multi-modal domain-aware conversation systems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.

Tulika Saha, Sriparna Saha, and Pushpak Bhattacharyya. 2020a. Towards sentiment aided dialogue policy learning for multi-intent conversations using hierarchical reinforcement learning. *PloS one*, 15(7):e0235367.

Tulika Saha, Sriparna Saha, and Pushpak Bhattacharyya. 2020b. Towards sentiment-aware multi-modal dialogue policy learning. *Cognitive Computation*, pages 1–15.

Weiyan Shi, Xuewei Wang, Yoo Jung Oh, Jingwen Zhang, Saurav Sahay, and Zhou Yu. 2020. Effects of persuasive dialogues: testing bot identities and inquiry strategies. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13.

Weiyan Shi and Zhou Yu. 2018. Sentiment adaptive end-to-end dialog systems. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1509–1519.

Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Yixuan Su, Yan Wang, Deng Cai, Simon Baker, Anna Korhonen, and Nigel Collier. 2021. Prototype-to-style: Dialogue generation with style-aware editing on retrieval memory. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:2152–2161.

Youzhi Tian, Weiyan Shi, Chen Li, and Zhou Yu. 2020. Understanding user resistance strategies in persuasive conversations. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4794–4798.

Abhisek Tiwari, Sriparna Saha, and Pushpak Bhattacharyya. 2022a. A knowledge infused context driven dialogue agent for disease diagnosis using hierarchical reinforcement learning. *Knowledge-Based Systems*, 242:108292.

Abhisek Tiwari, Tulika Saha, Sriparna Saha, Shubhashis Sengupta, Anutosh Maitra, Roshni Ramnani, and Pushpak Bhattacharyya. 2021a. A dynamic goal adapted task oriented dialogue agent. *Plos one*, 16(4):e0249030.

Abhisek Tiwari, Tulika Saha, Sriparna Saha, Shubhashis Sengupta, Anutosh Maitra, Roshni Ramnani, and Pushpak Bhattacharyya. 2021b. Multi-modal dialogue policy learning for dynamic and co-operative goal setting. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.

Abhisek Tiwari, Tulika Saha, Sriparna Saha, Shubhashis Sengupta, Anutosh Maitra, Roshni Ramnani, and Pushpak Bhattacharyya. 2022b. A persona aware persuasive dialogue policy for dynamic and co-operative goal setting. *Expert Systems with Applications*, 195:116303.

Hado Van Hasselt, Arthur Guez, and David Silver. 2016. Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.

Xuewei Wang, Weiyan Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. Persuasion for good: Towards a personalized persuasive dialogue system for social good. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5635–5649.

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

Anbang Xu, Zhe Liu, Yufan Guo, Vibha Sinha, and Rama Akkiraju. 2017. A new chatbot for customer service on social media. In *Proceedings of the 2017 CHI conference on human factors in computing systems*, pages 3506–3510.

Rui Yan. 2018. " chitty-chitty-chat bot": Deep learning for conversational ai. In *IJCAI*, volume 18, pages 5520–5526.

Zhirui Zhang, Xiujun Li, Jianfeng Gao, and Enhong Chen. 2019. Budgeted policy learning for task-oriented dialogue systems. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3742–3751.

A   Appendix

A.1   Dataset Details

Figure 5a shows word clouds of the curated conversational corpus. We also illustrated the word cloud of the corpus with persuasion strategy annotation in Figure 5b. We report meta data such as intent, slot and dialogue act lists in Table 10. Figure 6 and 7 show sentiment and persuasion strategy distribution across the corpus, respectively. We identified five image categories with 13 multi-modal attributes
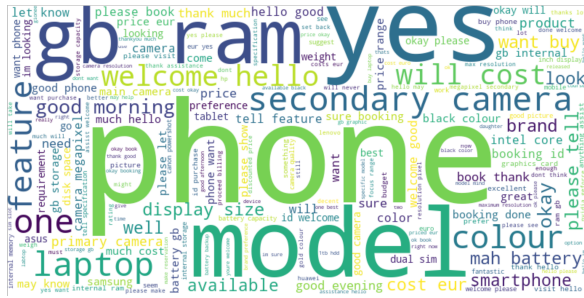
(a)   (b)

Figure 5: Word clouds for the curated PPMD corpus - a) user and agent conversations, b) user and agent conversations with persuasion strategy annotation

Intent	greet, specification, inform, request, persuasion, thanks, preq, done
Slot	model, brand, battery, ram, p_camera, s_camera, radio, display_size, status, sim, gps, os, color, internal_ram, weight, released_year, discount, released_month, price, phn_key, specifications, sp_done, features
Dialogue Act	greet, specification_request, specification_done, inform, request, result, recommend, persuade, re-persuade, goal_update, booking, close
Sentiment	positive, negative, neutral
Persuasion Strategy	Default, Credibility appeal, Logical appeal, Personal appeal, Emotional appeal, Persona based appeal

Table 10: Intent, slot and dialogue act list of the PPD dataset

Persuasion strategy	Example
Credibility appeal	It is a Nokia brand phone, which ensures its outstanding quality. Many other brand phones with the same quantity do not perform equally well for a long time. You should buy this phone without a second thought.
Logical appeal	You should buy this phone; it has lot of features such as a Radeon Pro 555X G2DDR5 (4 GB) graphic design with Intel Core i7 6 Core processor, 15.4 display size. Its rating is 4.1
Persona-based appeal	Sure, but I still highly recommend this phone to you because of its special features, particularly the gorgeous titan black color.
Emotional appeal	This phone will be a perfect gift for a photographer; it has all the features and specifications which are necessary for a photographer. Your girlfriend will love this for sure.
Personal appeal	This is a great phone and has received overwhelmingly positive reviews globally.

Table 11: Examples of different persuasion strategies

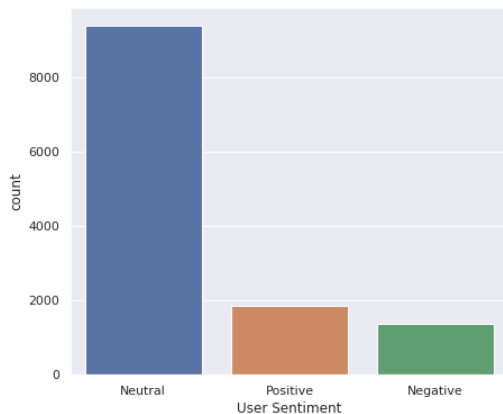of phone (Table 12) and tablet, which are hard to convey through text.

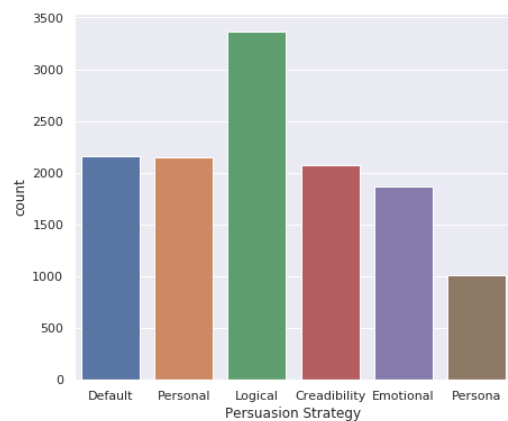

Figure 6: User sentiment distribution in the PPMD

Figure 7: User sentiment distribution in the PPMD

A.2 Implementation

The proposed methodology has been trained and evaluated on 80% and 20% of the complete dataset, respectively. Similar to other existing reinforce-
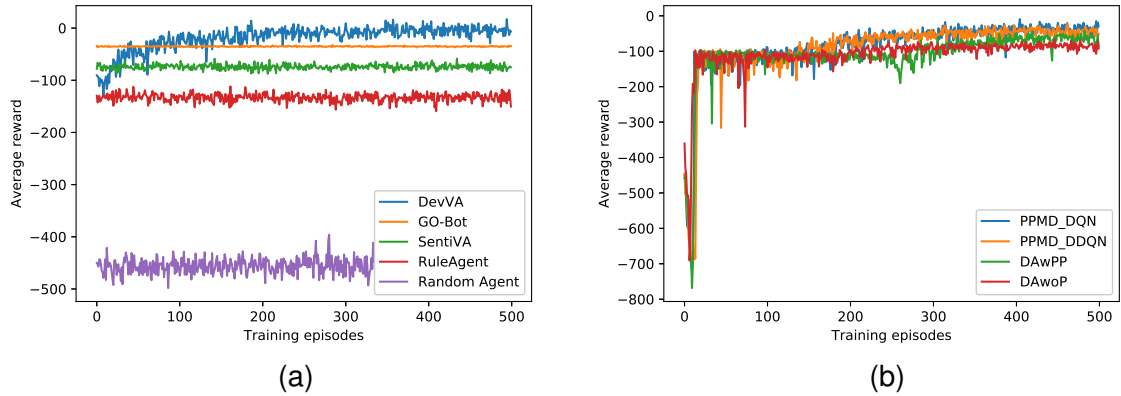
Figure 8: a. Avg. episodic reward of Random, Rule, SentiVA, Go-Bot and DevVA during training episodes, b. Avg. episodic reward of baselines and the proposed dialogue agent (PPMD) during training episodes

Category	Attributes	Number of samples
Color	Rose Gold, Black, Blue, Glacier White, Yellow, Silver	417
Style	Slide	555
Shape	Landscape	125
Type	Keypad	438
Brand	Apple, Samsung, MOTO, Huawei	326

Table 12: Different image categories and their multi-modal attributes

ment learning based dialogue agents, we have also utilized a user simulator for interacting with the dialogue agent. We developed an task-driven user simulator with reference the publicly available user simulator (Li et al., 2016). The model has been trained for 500 episodes, and each episode simulates 100 dialogues. The parameter values are as follows - TR_1: 3, TR_2: 5, TR_3: 2 , TR_4: 1, SR_1: 2, SR_2: 1, PR_1: 10 , PR_2: 20, PR_3: 5, learning rate :0.0001, p: 0.3, s: 0.3, $p_{success}$: 1, p_{fail}: 1. We report all the hperparameter values in Table 13. All the values are decided empirically.

Hyperparameter	Value
discount factor (γ)	0.9
batch size	32
train freq	100
learning rate	0.0001
Maximum dialogue length (N)	20
dqn_hidden_size	70
epsilon_initial	0.99
min_epsilon	0.01
epsilon_reduction_rate	0.0001

Table 13: Hyperparameter values

A.3 Analysis

Figure 8 illustrates the learning curve (in terms of episodic reward) of different baselines and the proposed model. We have shown the confusion matrix of the persuasion strategy classifier in Figure 9.
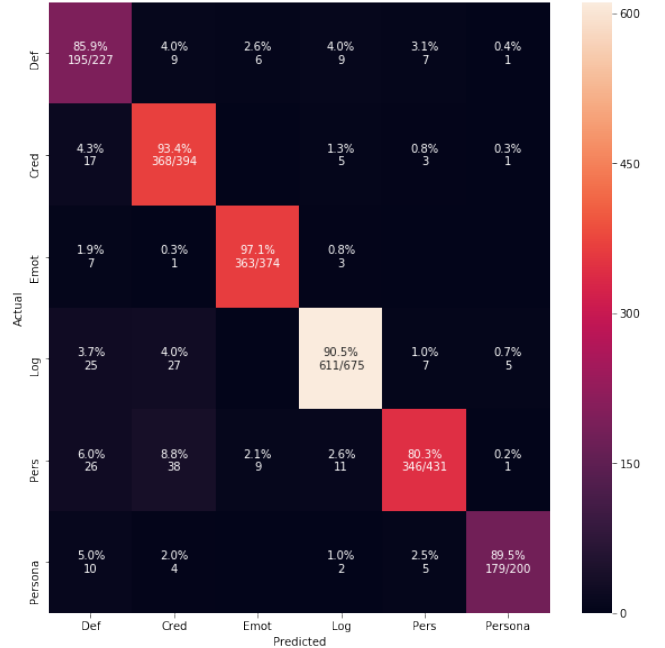

Figure 9: Confusion matrix of Persuasion strategy identifier