

LEARNING TIME-VARYING TURN-TAKING BEHAVIOR IN GROUP CONVERSATIONS

Madeline Navarro, Lisa O'Bryan, and Santiago Segarra

Rice University, Houston, TX, USA

ABSTRACT

We propose a flexible probabilistic model for *predicting turn-taking patterns in group conversations* based solely on *individual characteristics* and *past speaking behavior*. Many models of conversation dynamics cannot yield insights that generalize beyond a single group. Moreover, past works often aim to characterize speaking behavior through a universal formulation that may not be suitable for all groups. We thus develop a generalization of prior conversation models that predicts speaking turns among individuals in any group based on their individual characteristics, that is, personality traits, and prior speaking behavior. Importantly, our approach provides the novel ability to *learn* how speaking inclination varies based on when individuals last spoke. We apply our model to synthetic and real-world conversation data to verify the proposed approach and characterize real group interactions. Our results demonstrate that previous behavioral models may not always be realistic, motivating our data-driven yet theoretically grounded approach.

Index Terms— Conversational models, turn-taking, speaking behavior, probabilistic model, group interactions

1. INTRODUCTION

Understanding how individuals behave in group settings is valuable for both analyzing human behavior and improving communication-based tools. The ability to model group interactions can facilitate understanding of how they emerge from group composition variables, support the design of harmonious and effective teams based on team member compatibility, and enable the prediction of future group dynamics, informing potential interventions [1–3]. Moreover, insights into complex human interactions can contribute to building language models that are more realistic, understandable, or helpful [4].

While major developments in natural language processing have contributed to simulating realistic, text-based communication [5], there is still much to explore regarding how turn-taking behavior evolves in conversations and how team composition shapes these dynamics. Many past works that model conversation dynamics have aimed to capture universal conversation mechanics, which are generalizable but may not be suitable for all groups [6, 7]. Other existing approaches have limited investigation to the behavior of a single group and linked differences in speaking behaviors to specific individuals, limiting generalizability to other groups [8–10]. When it comes to the study of individual differences in speaking behaviors, traditional approaches generally involve interpretable yet simple tools such as those based on linear regression. For example,

studies have examined how personality traits, such as dominance, correlate with individual speaking time [11] or how mean team-level trait values, such as extraversion, correlate with the total number of utterances made by the team [12]. However, due to the interactive nature of conversational data [13], methods are required that take into account its dynamic and interdependent characteristics [14, 15]. Unfortunately, while data-driven tools demonstrate the ability to link individual characteristics to more complex behavioral patterns, the majority of such works implement black-box models, from which it is challenging to draw theoretical conclusions [16, 17].

We therefore propose a flexible parametric model for conversational turn-taking that predicts speaking patterns for groups of individuals. In particular, we generalize a well-founded statistical model that, at any point in a conversation, returns the probability that any individual will be the next speaker based on their personality traits and prior turn-taking behavior [6]. Our model is built by fitting behavioral parameters to observed conversation data among groups with known traits [18, 19]. We not only return scores reflecting the personalized speaking tendencies of any individual, but, differently from prior work [6, 7, 20], we also provide the novel ability to *learn* how speaking inclination changes based on how recently individuals last spoke. Thus, our model fitting process adapts to dynamic speaking patterns that may differ across group settings. This allows us to empirically evaluate if existing assumptions on how behavior varies over time align with real-world data [6, 8]. Moreover, beyond predictive ability, our use of machine learning constrained to theoretically grounded behavioral models yields interpretable characterizations of the relationship between personality and speaking behavior [7]. Our contributions are as follows.

1. We propose a flexible probabilistic model of turn-taking behavior for individuals in groups dependent upon their traits and conversation history, with the novel ability to learn how speaking inclination varies with the gap since individuals last spoke.
2. We verify our model by learning from both synthetic and real-world conversational data, evaluated by measuring our ability to predict conversations both globally and with an emphasis on rarer, more complex turn-taking patterns.
3. We compare our learned time-varying speaking behavior to existing assumptions, allowing us to investigate how real conversations evolve and thus assess the realism of common assumptions.

2. PROBLEM STATEMENT

We seek a model that predicts turn-taking patterns in a group conversation given observed traits of the group members. To this end, we consider a set of N individuals with measured traits collected in the vector $\mathbf{x} \in \mathbb{R}^N$, where $x_i \in \mathbb{R}$ denotes the trait of the i -th individual. These may correspond to any measured value associated with an individual, such as level of expertise on a group task; in this work, we focus on $\mathbf{x}$ representing personality traits. We emphasize our work with scalar-valued traits for simplicity, but exten-

Funding for this project was provided by the Army Research Institute for the Behavioral and Social Sciences (Grant W911NF-22-1-0226). The views, opinions, and/or findings contained in this article are those of the authors and shall not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documents. Emails: {nav, obryan, segarra}@rice.edu

sion to the multivariate setting is feasible and straightforward [7]. We further observe a sequence $\mathcal{C} := \{s(t)\}_{t=1}^T$ of T turns, where $s(t) \in [N] := \{1, 2, \dots, N\}$ represents which individual speaks at the t -th turn. We aim to obtain the likelihood that each individual will speak at the t -th turn given traits $\mathbf{x}$ and conversation history $\mathcal{H}_{t-1} := \{s(j)\}_{j=1}^{t-1} \subset \mathcal{C}$ up to turn $t-1$. To this end, we consider the following parametric approach to characterize speaking likelihoods based on a well-established conversational model [6]. At each turn t , the probability of member i speaking is $\mathbb{P}[s(t) = i] = u_i(t)/(1^\top \mathbf{u}(t))$, where $\mathcal{U} := \{\mathbf{u}(t)\}_{t=1}^T$ such that $\mathbf{u}(t) \in \mathbb{R}_+^N$ contains the time-varying speaking likelihoods of all N individuals. For a given member $i \in [N]$, we define

$$u_i(t) := \left[\pi_i + d_i w(\delta_i(t)) \right] \mathbb{I}\{\delta_i(t) > 1\}, \quad (1)$$

where $\mathbb{I}\{\cdot\}$ is the indicator function and $\pi, \mathbf{d} \in \mathbb{R}_+^N$ denote scores of speaking tendencies. The *inherent* tendency to speak for member i is π_i , while d_i reflects speaker *memory*, as d_i scales the time-dependent term in (1). The *proclivity* of a speaker $w : \mathbb{Z} \rightarrow \mathbb{R}_+$ represents how speaking inclination varies with the number of turns that have elapsed since an individual last spoke (previously termed the ‘memory function’ [7]), where $w(\delta) = 0$ for any $\delta \leq 0$, and $\delta_i(t) \in \mathbb{N}$ denotes the number of turns between t and the last turn in $\mathcal{H}_{t-1}$ spoken by i , with $\delta_i(t) = t - \max\{j : s(j) = i, j < t\}$ if member i has spoken at least once, otherwise, $\delta_i(t) = -\infty$. The first turn therefore depends only on the inherent scores $\mathbf{u}(1) = \pi$. Thus, larger values of w in (1) increase the chance that member i will speak, which may vary with how recently i last spoke. Observe that by the condition $\delta_i(t) > 1$, individuals cannot speak for consecutive turns, as we define a turn ending when another member speaks.

To discover how personality influences conversational behavior, we implement the model in (1) to predict speaking patterns from personality traits. Past works often assume $w(t) = e^{-b\delta_i(t)}$ for some $b > 0$, reflecting the greater tendency of recent speakers to speak again, with this likelihood decaying exponentially as the number of turns since they last spoke increases [6, 7, 20]. Differently, our representation in (1) allows for alternative patterns of time-varying behavior which can be learned directly from turn-taking data. Indeed, while $w(t) = e^{-b\delta_i(t)}$ provides a natural theoretical model, it may not reflect different patterns of speaking in all scenarios. Beyond modeling capabilities, existing work exploiting frameworks similar to (1) are typically employed to analyze the behavior of a given group based on observed conversation data, which cannot be generalized to unseen groups [6, 8, 20]. While progress has been made towards generalizability [7], adapting to time-varying patterns that may differ across group settings has not yet been explored.

3. METHODOLOGY

Our goal is to predict the likelihoods $\mathbf{u}(t)$ in (1) for any turn $t \in \mathbb{N}$ and any group of arbitrary size N given their traits $\mathbf{x} \in \mathbb{R}^N$ and conversation history $\mathcal{H}_{t-1} = \{s(j)\}_{j=1}^{t-1}$. To this end, we train models $f, g : \mathbb{R} \rightarrow \mathbb{R}_+$ and $\nu : \mathbb{Z} \rightarrow \mathbb{R}_+$ such that, for any $x \in \mathbb{R}, \delta \in \mathbb{Z}$,

$$a(x, \delta; \boldsymbol{\theta}) := \left[f(x; \boldsymbol{\theta}_1) + g(x; \boldsymbol{\theta}_2) \nu(\delta; \boldsymbol{\theta}_3) \right] \mathbb{I}\{\delta > 1\} \quad (2)$$

for learnable parameters $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \boldsymbol{\theta}_3)$. We may then let $\hat{u}_i(t) = a(x_i, \delta_i(t); \boldsymbol{\theta})$ to obtain speaking likelihood estimates $\hat{\mathcal{U}} = \{\hat{\mathbf{u}}(t)\}_{t=1}^T$, where $\hat{\pi}_i = f(x_i; \boldsymbol{\theta}_1)$ and $\hat{d}_i = g(x_i; \boldsymbol{\theta}_2)$ predict inherent and memory scores, respectively, and ν approximates speaking proclivity with $\nu(\delta; \boldsymbol{\theta}_3) = 0$ if $\delta \leq 0$. To obtain f, g ,

and ν , we consider G training groups, with traits $\mathbf{x}^{(g)} \in \mathbb{R}^{N_g}$ belonging to the g -th group, among whom we observe a conversation $\mathcal{C}^{(g)} = \{s^{(g)}(t)\}_{t=1}^{T_g}$. We learn the parameters $\boldsymbol{\theta}$ by maximizing the likelihood of each $\mathcal{C}^{(g)}$ given $\mathbf{x}^{(g)}$ via

$$\min_{\boldsymbol{\theta}} \sum_{g=1}^G \ell(\hat{\mathcal{U}}^{(g)}, \mathcal{C}^{(g)}) \quad \text{s.t.} \quad \hat{u}_i^{(g)}(t) = a(x_i^{(g)}, \delta_i^{(g)}(t); \boldsymbol{\theta}) \quad \forall i \in [N_g], t \in [T_g], g \in [G], \quad (3)$$

where, for any $\mathcal{U} = \{\mathbf{u}(t)\}_{t=1}^T$ and $\mathcal{C} \in [N]^T$, the loss function

$$\ell(\mathcal{U}, \mathcal{C}) := - \sum_{t=1}^T \log \frac{u_{s(t)}(t)}{1^\top \mathbf{u}(t)} \quad (4)$$

returns the negative log-likelihood of conversation $\mathcal{C}$ given $\mathcal{U}$ with $\mathbf{u}(t) \in \mathbb{R}_+^N$. While our approach supports several architectures, in this paper we let each of the components f, g , and ν be multi-layer perceptrons (MLPs) with a sigmoidal nonlinearity $\text{sig}(x) := (1 + e^{-x})^{-1}$ at the output since we aim to predict nonnegative values $\mathbf{u}(t)$. We solve (3) via block coordinate descent [21], where we alternately update $(\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ and $\boldsymbol{\theta}_3$ via gradient descent, precluding the need to compute gradients of the bilinear term in (2), which can lead to instabilities. Further details on both our architecture and simulations can be found in our GitHub repository¹.

By learning parameters of an interpretable and theoretically founded conversation model (2), we enjoy benefits distinct from other machine learning techniques, elaborated below. While sequential models such as recurrent neural networks (RNNs) are historical staples for time-series predictions [22], these do not allow generalizability across groups of different sizes. Indeed, tools that fix input dimensionality cannot describe turn-taking behavior across arbitrary groups. For example, Markov models have shown success in reflecting conversational patterns, albeit only for a single observed group [8]. In contrast, transformers allow arbitrary input and output dimensions and have been applied to create realistic sequences of data [23]. However, such architectures are also highly uninterpretable and therefore unlikely to soon be widely adopted for behavioral analyses [17]. In contrast, we follow similar paths as interdisciplinary approaches in science-informed machine learning, which train models restricted to well-understood constraints, such as theoretical rules or physical restrictions [24].

Solving (3) allows us to perform predictions on any set of individuals with known traits. First, given the trait x_i^* of any new individual, we may predict scores $\hat{\pi}_i^*$ and $\hat{d}_i^*$ via f and g , respectively, allowing us to characterize how behavior varies across personality traits. Second, for a set of individuals with traits $\mathbf{x}^* \in \mathbb{R}^{N^*}$ and conversation history $\mathcal{H}_{t-1}^*$, we predict the most likely next speaker as $\hat{s}^*(t) = \arg\max_i a(x_i^*, \delta_i^*(t); \boldsymbol{\theta})$ from (2). Third, we may generate realistic conversations from $\mathbf{x}^*$ by sampling speakers at each turn $s(t) \sim \hat{\mathbf{u}}(t)/(1^\top \hat{\mathbf{u}}(t))$ with $\hat{u}_i(t) = a(x_i, \delta_i(t); \boldsymbol{\theta})$. Thus, our approach advances capabilities of turn-taking analysis, sharing the goals of our past work [7]. However, ML-SPEAK in [7] did not learn a generic proclivity ν and instead assumed $w(\delta) = e^{-\delta/2}$, which we show may not be suited to all real-world conversations.

3.1. Evaluating predictive ability

To assess the ability of our model to accurately describe speaking tendencies, we measure the *test loss*, the loss for a conversation $\mathcal{C}^*$ among a new testing group given traits $\mathbf{x}^*$. To this end, we predict

¹Link to repository: <https://github.com/mn51/mlspeak-proclivity>




Fig. 1. (a) Testing loss for each baseline when $w = w_{\text{exp}}$. The left five boxplots are measured via ℓ in (4), the right five via ℓ_{turn} in (5). (b) Testing loss for each baseline when $w = w_{\text{sig}}$. (c) Testing loss for each baseline using real-world conversation data.

inherent and memory scores $\hat{\pi}^*$, $\hat{\mathbf{d}}^*$ from $\mathbf{x}^*$ and estimate speaking likelihoods $\hat{\mathcal{U}}^*$ as in (1) with the learned proclivity $\hat{w} = \nu$. Then, we compute the loss $\ell(\hat{\mathcal{U}}^*, \mathcal{C}^*)$ in (4), where lower values indicate that our model is well aligned with the dynamics in $\mathcal{C}^*$.

In addition to ℓ , we also propose an alternative loss, where each turn is weighted based on its *class*, determined by the pattern of turns taken prior $\mathcal{H}_t$. Following [6], we define the following $K := 4$ classes for a given turn t : (i) **Floor**: $s(t) = s(t-2)$; (ii) **Broken floor**: $s(t) \neq s(t-2)$, $s(t-1) = s(t-3)$; (iii) **Regain**: $s(t) = s(t-3) \neq s(t-1)$, $s(t-2) = s(t-4)$; (iv) **Nonfloor**: $\mathcal{H}_t$ does not satisfy (i), (ii), or (iii). Then, if T_k is the number of turns in class $k \in [K]$ and $\gamma(\mathcal{H}_t) := T/(KT_k)$ for turn t in class k , we define

$$\ell_{\text{turn}}(\mathcal{U}, \mathcal{C}) := - \sum_{t=1}^T \gamma(\mathcal{H}_t) \log \frac{u_{s(t)}(t)}{\mathbf{1}^\top \mathbf{u}(t)}. \quad (5)$$

Therefore, turns in classes that occur more frequently (larger T_k) are given smaller weights. As individuals who have recently spoken typically have a higher likelihood of speaking again, real-world conversations often contain a preponderance of floor turns [6, 8]. However, we wish to evaluate the ability of our model to predict a variety of patterns beyond those involving only two speakers. Thus, while the loss ℓ captures the overall likelihood of a conversation, ℓ_{turn} measures how well a model can predict more complex and realistic patterns by weighting the likelihoods of each turn class equitably.

4. RESULTS

We next apply our model (2) to synthetic and real-world data to predict speaking patterns. In particular, we compare the following approaches to estimate proclivity w and the scores $\pi, \mathbf{d} \in \mathbb{R}_+^N$ of any group with traits $\mathbf{x} \in \mathbb{R}^N$ and conversation $\mathcal{C} = \{s(t)\}_{t=1}^T$.

- **Proclivity (PRO)**: Our model (2) that predicts scores $\hat{\pi}_i^{\text{PRO}} = f(x_i; \theta_1)$, $\hat{d}_i^{\text{PRO}} = g(x_i; \theta_2)$ and learns proclivity $\hat{w}^{\text{PRO}}(\delta) = \nu(\delta; \theta_3)$, with $\theta = (\theta_1, \theta_2, \theta_3)$ learned via (3);
- **Exponential (EXP)** [7]: A similar model that predicts scores $\hat{\pi}_i^{\text{EXP}} = f(x_i; \theta_1)$, $\hat{d}_i^{\text{EXP}} = g(x_i; \theta_2)$ for a *fixed* proclivity $\hat{w}^{\text{EXP}}(\delta) = e^{-\delta/2}$, with θ_1, θ_2 learned via (3);
- **No Memory (NM)**: A memoryless model that predicts scores $\hat{\pi}_i^{\text{NM}} = 1$, $\hat{d}_i^{\text{NM}} = 0$ and ignores proclivity $\hat{w}^{\text{NM}}(\delta) = 0$,
- **High memory (HM)**: A high-memory model that predicts scores $\hat{\pi}_i^{\text{HM}} = 10^{-2}$, $\hat{d}_i^{\text{HM}} = 1$ with a fixed proclivity $\hat{w}^{\text{HM}}(\delta) = e^{-\delta/2}$.

Unlike our method **PRO**, **EXP** aligns with existing turn-taking analyses discussed in Section 2, which do not permit flexible modeling of how speaking inclination depends on the recency in which individuals last spoke [7]. The **NM** baseline predicts the next speaker

uniformly at random with no consideration of conversation history. Conversely, **HM** tends to select more recent speakers as the next predicted speaker, effectively ignoring inherent speaking tendencies.

4.1. Synthetically generated conversations

We first consider synthetic conversations to control the relationship between personality traits $\mathbf{x}$ and speaking likelihoods via $\pi, \mathbf{d}$. We generate 15 groups to be used for training and validation, each consisting of 5 members, where we sample traits $x_i^{(g)} \sim [0.1, 1]$ uniformly at random for each group $g \in [15]$. We then obtain the inherent and memory scores $\pi^{(g)}, \mathbf{d}^{(g)} \in \mathbb{R}_+^N$ through the mappings

$$\pi_i^{(g)} = \sqrt{x_i^{(g)}}, \quad d_i^{(g)} = \frac{15e}{2} \left(\frac{e^{-2(1.1-x_i^{(g)})} - e^{-2}}{e^{-0.2} - e^{-2}} + \frac{1}{3} \right), \quad (6)$$

chosen to yield different nonlinear relationships such that $\pi_i^{(g)} \in [0, 1]$ and $d_i^{(g)} \in [2.5e, 10e]$. Finally, we simulate a conversation per group based on $\pi^{(g)}, \mathbf{d}^{(g)}$ and a proclivity w of choice, yielding sequences $\mathcal{C}^{(g)} = \{s^{(g)}(t)\}_{t=1}^{800}$. We use 10 groups to train the parameters of the **PRO** and **EXP** models via (3), and the remaining 5 groups are set aside as validation data.

Furthermore, we generate 5 testing groups in the same manner, each with 5 speakers and a conversation of 800 turns, yielding traits $\mathbf{x}^{*(g)}$ and scores $\pi^{*(g)}, \mathbf{d}^{*(g)}$ for the g -th group, along with the conversation $\mathcal{C}^{*(g)}$ simulated with the same proclivity w . We evaluate each method by summing test losses across groups $\sum_g \ell(\hat{\mathcal{U}}^{*(g)}, \mathcal{C}^{*(g)})$ via (4) and $\sum_g \ell_{\text{turn}}(\hat{\mathcal{U}}^{*(g)}, \mathcal{C}^{*(g)})$ via (5), where $\hat{\mathcal{U}}^{*(g)}$ are the estimated likelihoods computed as in (1) from predictions $\hat{\pi}^{*(g)}, \hat{\mathbf{d}}^{*(g)}$, and $\hat{w}$.

We generate data for two types of proclivity,

$$w_{\text{exp}}(\delta) := e^{-\delta/2}, \quad w_{\text{sig}}(\delta) := 0.95 \text{sig}\left(10 - \frac{1}{2}\delta\right),$$

where $w_{\text{exp}}(\delta)$ aligns with assumptions in past works [6, 7, 20]. We further consider $w_{\text{sig}}(\delta)$, which also decreases as δ increases, yet as we will show, it renders the predictive task more challenging.

Exponentially decaying proclivity. First, we consider 10 independent trials of data generated with $w = w_{\text{exp}}$. The boxplots in Fig. 1a show the test losses summed over testing groups for ℓ in (4) and ℓ_{turn} in (5). Since we have ground truth scores $\pi^{*(g)}, \mathbf{d}^{*(g)}$ for each group and the true proclivity w_{exp} , we also compute test losses for the true scores $\mathcal{U}^{*(g)}$, denoted **True** in Fig. 1a. Note that Fig. 1a does not include **HM** for ℓ_{turn} as it attains significantly higher values (median approximately 1.56). We find that **PRO** and **EXP** yield lower ℓ than **True**, but they do not outperform **True** for the class-based loss ℓ_{turn} .




Fig. 2. (a) True versus learned proclivities for either a sigmoidal proclivity w_{sig} or exponentially decaying proclivity w_{exp} . (b) Learned proclivities for real-world conversation data.

This indicates that **PRO** and **EXP** tend to predict frequently occurring floor turns more accurately but may struggle with more complex patterns, as expected for data created using a proclivity w_{exp} that encourages floor turns. Indeed, observe that **HM**, which assumes the correct proclivity $\hat{w}^{\text{HM}} = w_{\text{exp}}$, attains a lower ℓ than **NM**, but the more naive **NM** obtains a superior loss ℓ_{turn} , implying that **HM** does not capture patterns accurately beyond floor turns.

Sigmoidal proclivity. We next consider the data sampled using a sigmoidal proclivity $w = w_{\text{sig}}$ for 10 trials, with the test losses for each method shown in Fig. 1b. Unlike for w_{exp} , we observe similar rankings between ℓ and ℓ_{turn} . First, note that **HM** with an inaccurate proclivity $\hat{w}^{\text{HM}}(\delta) = e^{-\delta/2}$ is not shown as it returns much higher losses ℓ and ℓ_{turn} than other methods (median values 1.69 and 1.55, respectively), where even the most naive approach **NM** outperforms **HM**. However, **EXP** implements the same incorrect proclivity as **HM** $\hat{w}^{\text{EXP}}(\delta) = e^{-\delta/2}$, but **EXP** far outperforms **HM** as it additionally learns individual scores $\hat{\pi}^{*(g)}, \hat{d}^{*(g)}$. Moreover, **PRO** exhibits superior loss to **EXP** for both ℓ and ℓ_{turn} , as **PRO** attains values much closer to those of **True**.

Visualizing predictions. We further visualize the proclivities for **PRO** and **EXP** along with the true functions for both $w \in \{w_{\text{exp}}, w_{\text{sig}}\}$ in Fig. 2a. While **True** and **EXP** do not learn speaking proclivity, we can illustrate the influence of the time-dependent term in (1) for each method by rescaling its proclivity by $\hat{\mathbb{E}}[d_i]/\hat{\mathbb{E}}[\pi_i]$, that is, the ratio of the empirical expected values of the memory and inherent scores, which are computed over a grid of equidistant traits $\mathbf{x} \in [0.1, 1]^{50}$. For w_{exp} we find that the rescaled proclivities of **PRO** and **EXP** are very similar, as suggested by the similarity of their predictive losses in 1a. However, for the more challenging w_{sig} , we find that **PRO** obtains a similarly decaying function to the true proclivity, albeit with a smaller effect for initial turns. Indeed, by our definition of w_{sig} , speakers may not experience a noticeable decay in speaking likelihood until approximately 10 turns elapse. Conversely, **EXP** learns to reduce dependence on time by returning a smaller ratio $\mathbb{E}[\hat{d}_i]/\mathbb{E}[\hat{\pi}_i]$, hence its lower rescaled proclivity in Fig 2 for all $\delta \geq 2$. While this mitigates the effect of its incorrectly assumed proclivity $\hat{w}^{\text{EXP}}(\delta) = e^{-\delta/2} \neq w_{\text{sig}}$, it also renders a worse predictive ability for **EXP**, as observed in Fig. 1b.

4.2. Real-world Data

Finally, we repeat the above experiments for real-world conversational data, described in greater detail in [7]. In brief, we observe a set of $G = 20$ teams of students among whom we observe conversations over Zoom, in which they discussed course projects. Further-

more, we also have measurements of each student’s level of extraversion $\{\mathbf{x}^{(g)}\}_{g=1}^{20}$, a personality trait that has a well-studied tendency to correlate with interactive behavior [12, 25]. We create 20 trials by first randomly selecting 12 training, 4 validation, and 4 testing groups for the first trial. Then, each subsequent trial corresponds to rotating group indices by one, such that every group is present or absent in training, validation, or testing sets for at least one trial. As we no longer have ground truth likelihoods $\mathcal{U}^*$, we compare test losses **PRO**, **EXP**, **NM**, and **HM** using predicted scores $\hat{\pi}^{*(g)}, \hat{d}^{*(g)}$ of each testing group along with the proclivity $\hat{w}$ for each method.

Boxplot comparisons of test losses ℓ and ℓ_{turn} are shown in Fig. 1c. We find that **PRO** is competitive with **EXP** and **HM** for the unweighted loss ℓ . As mentioned previously, real-world conversations often contain large numbers of floor turns. Thus, methods such as **EXP** and **HM** can obtain low values of ℓ since they assume an exponentially decaying proclivity w_{exp} , which tends to encourage predicting floor turns. Aligning with our intuition, **NM**, which does not employ time-varying predictions, yields the highest ℓ for the real-world conversations. However, when we observe ℓ_{turn} , for which lower values indicate greater ability to predict more complex turn-taking patterns, all methods **EXP**, **NM**, and **HM** attain significantly worse values than **PRO**. Thus, our approach **PRO**, which learns time-varying proclivity, indeed exhibits flexibility when adapting to real-world conversational data, as we are able to demonstrate superior predictive ability for non-trivial turn-taking patterns.

Moreover, we visualize proclivities for the real-world data in Fig. 2b, with $\hat{w}^{\text{EXP}}$ and $\hat{w}^{\text{PRO}}$ scaled by inherent and memory scores as described previously. As we do not have ground truth proclivities, we instead plot w_{exp} rescaled to compare its shape to $\hat{w}^{\text{PRO}}$, as w_{exp} is a common choice for proclivity in past works [6, 7, 20]. First, we note that **EXP** reduces the effect of $\hat{w}^{\text{EXP}}$, as Fig. 2b shows lower values of its scaled proclivity. This implies that temporal effects are more subdued, potentially indicating that **EXP** employs an incorrect proclivity, which we also observed for the synthetic case $w = w_{\text{sig}}$. For our method **PRO**, our time-dependent likelihoods decay more slowly than w_{exp} , where $\hat{w}^{\text{PRO}}$ yields much higher values in Fig 2b than both w_{exp} and $\hat{w}^{\text{EXP}}$, especially when δ is approximately between 5 and 10. Thus, we find that these real-world conversations did not exhibit the typically assumed proclivity w_{exp} . Overall, we find that our **PRO** model in (2) is able to identify time-dependent speaking behavior for real group interactions shown in Fig. 2b, and we demonstrate its superior predictive ability for realistic, non-trivial speaking patterns in Fig. 1c in comparison with other methods.

5. CONCLUSION

In this work, we presented a flexible computational model that predicts the likelihood that each group member will speak at a given speaking turn in a conversation based on their personality traits and conversation history. Moreover, we introduced a novel approach to learn how speaking proclivity depends on how recently individuals have spoken, which allows our model to adapt to group settings with unique conversation dynamics. In addition, we applied our model to real-world conversation data and compared the learned time-varying proclivities with previous assumptions of how speaking behavior evolves. We found that the real-world group conversations that we observed exhibited different behavior than is commonly assumed, demonstrating the need to consider variation in group dynamics and also develop flexible tools to model them. In future work, we will consider proclivity models more amenable to time-series predictions.

6. REFERENCES

- [1] Suzanne T Bell, Shanique G Brown, Anthony Colaneri, and Neal Outland, “Team composition and the ABCs of teamwork,” *American Psychologist*, vol. 73, no. 4, pp. 349–362, 2018.
- [2] Lisa O’Bryan, Margaret Beier, and Eduardo Salas, “How approaches to animal swarm intelligence can improve the study of collective intelligence in human teams,” *J. of Intell.*, vol. 8, no. 1, pp. 9, 2020.
- [3] Steve W J Kozlowski and Georgia T Chao, “Unpacking team process dynamics and emergent phenomena: Challenges, conceptual advances, and innovative methods,” *American Psychologist*, vol. 73, no. 4, pp. 576, 2018.
- [4] Xuhui Zhou, Zhe Su, Tiwalayo Eisape, Hyunwoo Kim, and Maarten Sap, “Is this the real life? Is this just fantasy? The misleading success of simulating social interactions with LLMs,” in *Proc. of the Conf. on Empirical Methods in Natural Language Process*. 2024, pp. 21692–21714, ACL.
- [5] Hyunwoo Kim, Jack Hessel, Liwei Jiang, Peter West, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Bras, Malihe Alikhani, Gunhee Kim, Maarten Sap, and Yejin Choi, “SODA: Million-scale dialogue distillation with social commonsense contextualization,” in *Proc. of the Conf. on Empirical Methods in Natural Language Process*. 2023, pp. 12930–12949, ACL.
- [6] Garold Stasser and Laurie A Taylor, “Speaking turns in face-to-face discussions,” *J. of Person. and Soc. Psych.*, vol. 60, no. 5, pp. 675–684, 1991.
- [7] Lisa O’Bryan, Madeline Navarro, Juan Segundo Hevia, and Santiago Segarra, “ML-SPEAK: A theory-guided machine learning method for studying and predicting conversational turn-taking patterns,” *J. of Person. and Soc. Psych.*, 2025.
- [8] Kevin C Parker, “Speaking turns in small group interaction: A context-sensitive event sequence model,” *J. of Person. and Soc. Psych.*, vol. 54, no. 6, pp. 965–971, 1988.
- [9] Sumit Basu, Tanzeem Choudhury, and Brian Clarkson, “Learning human interactions with the influence model,” Tech. Rep. MIT-TR-539, MIT Media Laboratory, 2001.
- [10] Emiliano Padilha and Jean Carletta, “A simulation of small group discussion,” in *Wrkshp. on the Semantics and Pragmatics of Dialogue*, 2002, pp. 117–124.
- [11] Glenn E Littlepage, Greg W Schmidt, Eric W Whisler, and Alan G Frost, “An input-process-output analysis of influence and performance in problem-solving groups,” *J. of Person. and Soc. Psych.*, vol. 69, no. 5, pp. 877, 1995.
- [12] Gretchen A Macht, David A Nembhard, Jung Hyup Kim, and Ling Rothrock, “Structural models of extraversion, communication, and team performance,” *Intl. J. of Industrial Ergonomics*, vol. 44, no. 1, pp. 82–91, 2014.
- [13] Harvey Sacks, Emanuel A Schegloff, and Gail Jefferson, “A simplest systematics for the organization of turn-taking for conversation,” *Language*, vol. 50, no. 4, pp. 696–735, 1974.
- [14] Goran Kuljanin, Michael T. Braun, James A. Grand, Jeffrey D. Olenick, Georgia T. Chao, and Steve W.J. Kozlowski, “Advancing organizational science with computational process theories,” *The Leadership Quarterly*, vol. 35, no. 4, pp. 101797, 2024.
- [15] Wiebke Bleidorn and Christopher James Hopwood, “Using machine learning to advance personality assessment and theory,” *Person. and Soc. Psych. Review*, vol. 23, no. 2, pp. 190–203, 2019.
- [16] Tal Yarkoni and Jacob Westfall, “Choosing prediction over explanation in psychology: Lessons from machine learning,” *Perspectives on Psychological Science*, vol. 12, no. 6, pp. 1100–1122, 2017.
- [17] Clemens Stachl, Florian Pargent, Sven Hilbert, Gabriella M Harari, Ramona Schoedel, Sumer Vaid, Samuel D Gosling, and Markus Böhner, “Personality research and assessment in the era of machine learning,” *European J. of Person.*, vol. 34, no. 5, pp. 613–631, 2020.
- [18] Charles E. Kimble and Jeffrey I. Musgrove, “Dominance in arguing mixed-sex dyads: Visual dominance patterns, talking time, and speech loudness,” *J. of Res. in Person.*, vol. 22, no. 1, pp. 1–16, 1988.
- [19] Ronen Cuperman and William Ickes, “Big Five predictors of behavior and perceptions in initial dyadic interactions: Personality similarity helps extraverts and introverts, but hurts “disagreeables”,” *J. of Person. and Soc. Psych.*, vol. 97, no. 4, pp. 667, 2009.
- [20] Lisa O’Bryan, Santiago Segarra, Jensine Paoletti, Stephanie Zajac, Margaret E Beier, Ashutosh Sabharwal, Matthew Wetergreen, and Eduardo Salas, “Conversational turn-taking as a stochastic process on networks,” in *Asilomar Conf. Signals, Syst., and Computers*, 2022, pp. 1243–1247.
- [21] Dimitri Bertsekas, *Nonlinear Programming*, Athena Scientific, 1999.
- [22] Ilya Sutskever, Oriol Vinyals, and Quoc V Le, “Sequence to sequence learning with neural networks,” in *Advances in Neural Info. Process. Syst.* 2014, vol. 27, Curran Associates, Inc.
- [23] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *Advances in Neural Info. Process. Syst.* 2017, vol. 30, Curran Associates, Inc.
- [24] Ali Ghaseminejad and Venkatesh Uddameri, “Physics-inspired integrated space-time artificial neural networks for regional groundwater flow modeling,” *Hydrology and Earth Syst. Sciences*, vol. 24, no. 12, pp. 5759–5779, 2020.
- [25] Shui-Kwan Leung and Michael Harris Bond, “Interpersonal communication and personality: Self and other perspectives,” *Asian J. of Soc. Psych.*, vol. 4, no. 1, pp. 69–86, 2001.