

Discovering User Groups for Natural Language Generation

Nikos Engonopoulos and Christoph Teichmann and Alexander Koller
Saarland University

{nikos|cteichmann|koller}@coli.uni-saarland.de

Abstract

We present a model which predicts how individual users of a dialog system understand and produce utterances based on user groups. In contrast to previous work, these user groups are not specified beforehand, but learned in training. We evaluate on two referring expression (RE) generation tasks; our experiments show that our model can identify user groups and learn how to most effectively talk to them, and can dynamically assign unseen users to the correct groups as they interact with the system.

1 Introduction

People vary widely both in their linguistic preferences when producing language and in their ability to understand specific natural-language expressions, depending on what they know about the domain, their age and cognitive capacity, and many other factors. It has long been recognized that effective NLG systems should therefore *adapt* to the current user, in order to generate language which works well for them. This adaptation needs to address all levels of the NLG pipeline, including discourse planning (Paris, 1988), sentence planning (Walker et al., 2007), and RE generation (Janarthanam and Lemon, 2014), and depends on many features of the user, including level of expertise and language proficiency, age, and gender.

Existing techniques for adapting the output of an NLG system have shortcomings which limit their practical usefulness. Some systems need user-specific information in training (Ferreira and Paraboni, 2014) and therefore cannot generalize to unseen users. Other systems assume that each user in the training data is annotated with their group, which allows them to learn a model from the data of each group. However, hand-designed user groups

may not reflect the true variability of the data, and may therefore inhibit the system’s ability to flexibly adapt to new users.

In this paper, we present a user adaptation model for NLG systems which induces user groups from training data in which these groups were not annotated. At training time, we probabilistically assign users to groups and learn the language preferences for each group. At evaluation time, we assume that our system has a chance to interact with each new user repeatedly – e.g., in the context of a dialogue system. It will then calculate an increasingly accurate estimate of the user’s group membership based on observable behavior, and use it to generate utterances that are suitable to the user’s true group.

We evaluate our model on two tasks involving the generation of referring expressions (RE). First, we predict the use of spatial relations in human-like REs in the GRE3D domain (Viethen and Dale, 2010) using a log-linear production model in the spirit of Ferreira and Paraboni (2014). Second, we predict the comprehension of generated REs, in a synthetic dataset based on data from the GIVE Challenge domain (Striegnitz et al., 2011) with the log-linear comprehension model of Engonopoulos et al. (2013). In both cases, we show that our model discovers user groups in the training data and infers the group of unseen users with high confidence after only a few interactions during testing. In the GRE3D domain, our system outperformed a strong baseline which used demographic information for the users.

2 Related Work

Differences between individual users have a substantial impact on language comprehension. Factors that play a role include level of expertise and spatial ability (Benyon and Murray, 1993); age (Häuser et al., 2017); gender (Dräger and Koller,

2012); or language proficiency (Koller et al., 2010).

Individual differences are also reflected in the way people produce language. Viethen and Dale (2008) present a corpus study of human-produced REs (GRE3D3) for simple visual scenes, where they note two clearly distinguishable groups of speakers, one that always uses a spatial relation and one that never does. Ferreira and Paraboni (2014) show that a model using speaker-specific information outperforms a generic model in predicting the attributes used by a speaker when producing an RE. However, their system needs to have seen the particular speaker in training, while our system can dynamically adapt to unseen users. Ferreira and Paraboni (2017) also demonstrate that splitting speakers in predefined groups and training each group separately improves the human likeness of REs compared to training individual user models.

The ability to adapt to the comprehension and production preferences of a user is especially important in the context of a dialog system, where there are multiple chances of interacting with the same user. Some methods adapt to dialog system users by explicitly modeling the users’ knowledge state. An early example is Paris (1988); she selects a discourse plan for a user, depending on their level of domain knowledge ranging between novice and expert, but provides no mechanism for inferring the group to which the user belongs. Rosenblum and Moore (1993) try to infer what knowledge a user possesses during dialogue, based on the questions they ask. Janarthanam and Lemon (2014) adapt to unseen users by using reinforcement learning with simulated users to make a system able to adjust to the level of the user’s knowledge. They use five predefined groups from which they generate the simulated users’ behavior, but do not assign real users to these groups. Our system makes no assumptions about the user’s knowledge and does not need to train with simulated users, or use any kind of information-seeking moves; we instead rely on the groups that are discovered in training and dynamically assign new, unseen users, based only on their observable behavior in the dialog.

Another example of a user-adapting dialog component is SPaRKY (Walker et al., 2007), a trainable sentence planner that can tailor sentence plans to individual users’ preferences. This requires training on separate data for each user; in contrast to this, we leverage the similarities between users and can take advantage of the full training data.

3 Log-linear models for NLG in dialog

We start with a basic model of the way in which people produce and comprehend language. In order to generalize over production and comprehension, we will simply say that a human language user exhibits a certain *behavior* b among a range of possible behaviors, in response to a *stimulus* s . The behavior of a speaker is the utterance b they produce in order to achieve a communicative goal s ; the behavior of a listener is the meaning b which they assign to the utterance s they hear.

Given this terminology, we define a basic log-linear model (Berger et al., 1996) of language use as follows:

$$P(b|s; \rho) = \frac{\exp(\rho \cdot \phi(b, s))}{\sum_{b'} \exp(\rho \cdot \phi(b', s))} \quad (1)$$

where ρ is a real-valued parameter vector of length n and $\phi(b, s)$ is a vector of real-valued *feature functions* f_1, \dots, f_n over behaviors and stimuli. The parameters can be trained by maximum-likelihood estimation from a corpus of observations (b, s) . In addition to maximum-likelihood training it is possible to include some prior probability distribution, which expresses our belief about the probability of any parameter vector and which is generally used for regularization. The latter case is referred to as *a posteriori* training, which selects the value of ρ that maximizes the product of the parameter probability and the probability of the data.

In this paper, we focus on the use of such models in the context of the NLG module of a dialogue system, and more specifically on the generation of referring expressions (REs). Using (1) as a *comprehension* model, Engonopoulos et al. (2013) developed an RE generation model in which the stimulus $s = (r, c)$ consists of an RE r and a visual context c of the GIVE Challenge (Striegnitz et al., 2011), as illustrated in Fig. 1. The behavior is the object b in the visual scene to which the user will resolve the RE. Thus for instance, when we consider the RE $r =$ “the blue button” in the context of Fig. 1, the log-linear model may assign a higher probability to the button on the right than to the one in the background. Engonopoulos and Koller (2014) develop an algorithm for generating the RE r which maximizes $P(b^*|s; \rho)$, where b^* is the intended referent in this setting.

Conversely, log-linear models can also be used to directly capture how a human speaker would refer to an object in a given scene. In this case, the stimulus $s = (a, c)$ consists of the target object a and

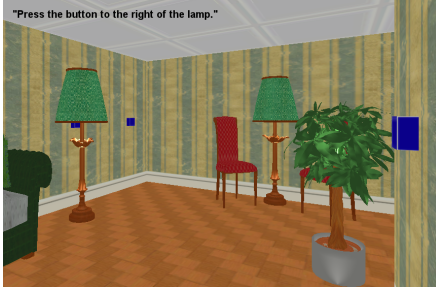


Figure 1: A visual scene and a system-generated instruction from the GIVE challenge.

the visual context c , and the behavior b is the RE. We follow Ferreira and Paraboni (2014) in training individual models for the different attributes which can be used in the RE (e.g., that a is a button; that it is blue; that the RE contains a binary relation such as “to the right of”), such that we can simply represent b as a binary choice $b \in \{1, -1\}$ between whether a particular attribute should be used in the RE or not. We can then implement an analog of Ferreira’s model in terms of (1) by using feature functions $\phi(b, a, c) = b \cdot \phi'(a, c)$, where $\phi'(a, c)$ corresponds to their *context* features, which do not capture any speaker-specific information.

4 Log-linear models with user groups

As discussed above, a user-agnostic model such as (1) does not do justice to the variability of language comprehension and production across different speakers and listeners. We will therefore extend it to a model which distinguishes different *user groups*. We will not try to model why¹ users behave differently. Instead our model sorts users into groups simply based on the way in which they respond to stimuli, in the sense of Section 3, and implements this by giving each group g its own parameter vector $\rho^{(g)}$. As a theoretical example, Group 1 might contain users who reliably comprehend REs which use colors (“the green button”), whereas Group 2 might contain users who more easily understand relational REs (“the button next to the lamp”). These groups are then discovered at training time.

When our trained NLG system starts interacting with an unseen user u , it will infer the group to which u belongs based on u ’s observed responses to previous stimuli. Thus as the dialogue with u unfolds, the system will have an increasingly pre-

¹E.g., in the sense of explicitly modeling sociolects or the difference between novice system users vs. experts.

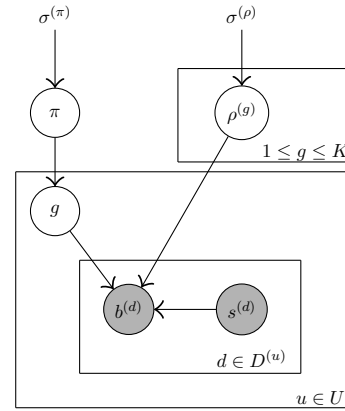


Figure 2: Plate diagram for the user group model.

cise estimate of the group to which u belongs, and will thus be able to generate language which is increasingly well-tailored to this particular user.

4.1 Generative story

We assume training data $D = \{(b_i, s_i, u_i)\}_i$ which contains stimuli s_i together with the behaviors b_i which the users u_i exhibited in response to s_i . We write $D^{(u)} = \{(b_1^u, s_1^u), \dots, (b_N^u, s_N^u)\}$ for the data points for each user u .

The generative story we use is illustrated in Fig. 2; observable variables are shaded gray, unobserved variables and parameters to be set in training are shaded white and externally set hyperparameters have no circle around them. Arrows indicate which variables and parameters influence the probability distribution of other variables.

We assume that each user belongs to a group $g \in \{1, \dots, K\}$, where the number K of groups is fixed beforehand based on, e.g., held out data. A group g is assigned to u at random from the distribution

$$P(g|\pi) = \frac{\exp(\pi_g)}{\sum_{g'=1}^K \exp(\pi_{g'})} \quad (2)$$

Here $\pi \in \mathbb{R}^K$ is a vector of weights, which defines how probable each group is a-priori.

We replace the single parameter vector ρ of (1) with group-specific parameters vectors $\rho^{(g)}$, thus obtaining a potentially different log-linear model $P(b|s; \rho^{(g)})$ for each group. After assigning a group, our model generates responses b_1^u, \dots, b_N^u at random from $P(b|s; \rho^{(g)})$, based on the group specific parameter vector and the stimuli s_1^u, \dots, s_N^u . This accounts for the generation of the data.

We model the parameter vectors $\pi \in \mathbb{R}^K$, and $\rho^{(g)} \in \mathbb{R}^n$ for every $1 \leq g \leq K$ as drawn from

$$P(D; \theta) = \left(\prod_{u \in U} \sum_{g=1}^K P(g|\pi) \cdot \prod_{d \in D^{(u)}} P(b_d|s_d; \rho^{(g)}) \right) \cdot \left(\mathcal{N}(\pi|0, \sigma^{(\pi)}) \cdot \prod_{g=1}^K \mathcal{N}(\rho^{(g)}|0, \sigma^{(\rho)}) \right) \quad (3)$$

$$\mathcal{L}(\theta) = \sum_{u \in U} \log \sum_{g=1}^K P(g|\pi) \cdot \prod_{d \in D^{(u)}} P(b_d|s_d; \rho^{(g)}) \quad (4)$$

$$\mathcal{AL}(\theta) = \sum_{u \in U} \sum_{g=1}^K \left(P(g|D^{(u)}; \theta_{(i-1)}) \cdot \left(\log P(g|\pi) + \sum_{d \in D_u} \log P(b_d|s_d; \rho^{(g)}) \right) \right) \quad (5)$$

normal distributions $\mathcal{N}(0, \sigma^{(\pi)})$, and $\mathcal{N}(0, \sigma^{(\rho)})$, which are centered at 0 with externally given variances and no covariance between parameters. This has the effect of making parameter choices close to zero more probable. Consequently, our models are unlikely to contain large weights for features that only occurred a few times or which are only helpful for a few examples. This should reduce the risk of overfitting the training set.

The equation for the full probability of the data and a specific parameter setting is given in (3). The left bracket contains the likelihood of the data, while the right bracket contains the prior probability of the parameters.

4.2 Predicting user behavior

Once we have set values $\theta = (\pi, \rho^{(1)}, \dots, \rho^{(K)})$ for all the parameters, we want to predict what behavior b a user u will exhibit in response to a stimulus s . If we encounter a completely new user u , the prior user group distribution from (2) gives the probability that this user belongs to each group. We combine this with the group-specific log-linear behavior models to obtain the distribution:

$$P(b|s; \theta) = \sum_{g=1}^K P(b|s; \rho^{(g)}) \cdot P(g|\pi) \quad (6)$$

Thus, we have a group-aware replacement for (1).

Furthermore, in the interactive setting of a dialogue system, we may have multiple opportunities to interact with the same user u . We can then develop a more precise estimate of u 's group based on their responses to previous stimuli. Say that we have made the previous observations $D^{(u)} = \{\langle s_1, b_1 \rangle, \dots, \langle s_N, b_N \rangle\}$ for user u . Then we can use Bayes' theorem to calculate a *posterior* estimate for u 's group membership:

$$P(g|D^{(u)}; \theta) \propto P(D^{(u)}|\rho^{(g)}) \cdot P(g|\pi) \quad (7)$$

This posterior balances whether a group is likely in general against whether members of that group behave as u does. We can use $P_u(g) = P(g|D^{(u)}; \theta)$ as our new estimate for the group membership probabilities for u and replace (6) with:

$$P(b|s, D^{(u)}; \theta) = \sum_{g=1}^K P(b|s; \rho^{(g)}) \cdot P_u(g) \quad (8)$$

for the next interaction with u .

An NLG system can therefore adapt to each new user over time. Before the first interaction with u , it has no specific information about u and models u 's behavior based on (6). As the system interacts with u repeatedly, it collects observations $D^{(u)}$ about u 's behavior. This allows it to calculate an increasingly accurate posterior $P_u(g) = P(g|D^{(u)}; \theta)$ of u 's group membership, and thus generate utterances which are more and more suitable to u using (8).

5 Training

So far we have not discussed how to find settings for the parameters $\theta = \pi, \rho^{(1)}, \dots, \rho^{(K)}$, which define our probability model. The key challenge for training is the fact that we want to be able to train while treating the assignment of users to groups as unobserved.

We will use a maximum *a posteriori* estimate for θ , i.e., the setting which maximizes (3) when D is our training set. We will first discuss how to pick parameters to maximize only the left part of (3), i.e., the data likelihood, since this is the part that involves unobserved variables. We will then discuss handling the parameter prior in section 5.2.

5.1 Expectation Maximization

Gradient descent based methods (Nocedal and Wright, 2006) exist for finding the parameter settings which maximize the likelihood for log-linear

models, under the conditions that all relevant variables are observed in the training data. If group assignments were given, gradient computations, and therefore gradient based maximization, would be straightforward for our model. One algorithm specifically designed to solve maximization problems with unknown variables by reducing them to the case where all variables are observed, is the expectation maximization (EM) algorithm (Neal and Hinton, 1999). Instead of maximizing the data likelihood from (3) directly, EM equivalently maximizes the log-likelihood, given in (4). It helps us deal with unobserved variables by introducing “pseudo-observations” based on the expected frequency of the unobserved variables.

EM is an iterative algorithm which produces a sequence of parameter settings $\theta_{(1)}, \dots, \theta_{(n)}$. Each will achieve a larger value for (4). Each new setting is generated in two steps: (1) an lower bound on the log-likelihood is generate and (2) the new parameter setting is found by optimizing this lower bound. To find the lower bound we compute the probability for every possible value the unobserved variables could have had, based on the observed variables and the parameter setting $\theta_{(i-1)}$ from the last iteration step. Then the lower bound essentially assumes that each assignment was seen with a frequency equal to these probabilities - these are the “pseudo-observations”.

In our model the unobserved variables are the assignments of users to groups. The probability of seeing each user u assigned to a group, given all the data $D^{(u)}$ and the model parameters from the last iteration $\theta_{(i-1)}$, is simply the posterior group membership probability $P(g|D^{(u)}; \theta_{(i-1)})$. The lower bound is then given by (5). This is the sum of the log probabilities of the data points under each group model, weighted by $P(g|D^{(u)}; \theta_{(i-1)})$. We can now use gradient descent techniques to optimize this lower bound.

5.1.1 Maximizing the Lower Bound

To fully implement EM we need a way to maximize (5). This can be achieved with gradient based methods such as L-BFGS (Nocedal and Wright, 2006). Here the gradient refers to the vector of all partial derivatives of the function with respect to each dimension of θ . We therefore need to calculate these partial derivatives.

There are existing implementations of the gradient computations our base model such as in Engonopoulos et al. (2013). The gradients of (5)

for each of the $\rho^{(g)}$ is simply the gradient for the base model on each datapoint d weighted by $P(g|D^{(u)}; \theta_{(i-1)})$ if $d \in D_u$, i.e., the probability that the user u from which the datapoint originates belongs to group g . We can therefore compute the gradients needed for each $\rho^{(g)}$ by using implementations developed for the base model.

We also need gradients for the parameters in π , which are only used in our extended model. We can use the rules for computing derivatives to find, for each dimension g :

$$\frac{\partial \mathcal{L}(\theta)}{\partial \pi_g} = \sum_{u \in U} P_u(g) - \frac{\exp(\pi_g)}{\sum_{g'=1}^K \exp(\pi_{g'})}$$

where $P_u(g) = P(g|D^{(u)}; \theta_{(i-1)})$. Using these gradients we can use L-BFGS to maximize the lower bound and implement the EM iteration.

5.2 Handling the Parameter Prior

So far we have discussed maximization only for the likelihood without accounting for the prior probabilities for every parameter. To obtain our full training objective we add the log of the right hand side of (3):

$$\log \left(\mathcal{N}(\pi|0, \sigma^{(\pi)}) \cdot \prod_{g=1}^K \mathcal{N}(\rho^{(g)}|0, \sigma^{(\rho)}) \right)$$

i.e., the parameter prior, to (4) and (5). The gradient contribution from these priors can be computed with standard techniques.

5.3 Training Iteration

We can now implement an EM loop, which maximizes (3) as follows: we randomly pick an initial value $\theta_{(0)}$ for all parameters. Then we repeatedly compute the $P(g|D^{(u)}; \theta_{(i-1)})$ values and maximize the lower bound using L-BFGS to find $\theta_{(i)}$. This EM iteration is guaranteed to eventually converge towards a local optimum of our objective function. Once change in the objective falls below a pre-defined threshold, we keep the final θ setting.

For our implementation we make a small improvement to the approach: L-BFGS is itself an iterative algorithm and instead of running it until convergence every time we need to find a new $\theta_{(i)}$, we only let it take a few steps. Even if we just took a single L-BFGS step in each iteration, we would still obtain a correct algorithm (Neal and

Hinton, 1999) and this has the advantage that we do not spend time trying to find a $\theta_{(i)}$ which is a good fit for the likely poor group assignments $P(g|D^{(u)}; \theta_{(i-1)})$ we obtain from early parameter estimates.

6 Evaluation

Our model can be used in any component of a dialog system for which a prediction of the user’s behavior is needed. In this work, we evaluate it in two NLG-related prediction tasks: RE production and RE comprehension. In both cases we evaluate the ability of our model to predict the user’s behavior given a stimulus. We expect our user-group model to gradually improve its prediction accuracy compared to a generic baseline without user groups as it sees more observations from a given user.

In all experiments described below we set the prior variances $\sigma_\gamma = 1.0$ and $\sigma_\pi = 0.3$ after trying out values between 0.1 and 10 on the training data of the comprehension experiment.

6.1 RE production

Task The task of RE generation can be split in two steps: *attribute selection*, the selection of the visual attributes to be used in the RE such as color, size, relation to other objects and *surface realization*, the generation of a full natural language expression. We focus here on attribute selection: given a visual scene and a target object, we want to predict the set of attributes of the target object that a human speaker would use in order to describe it. Here we treat attribute selection in terms of individual classification decisions on whether to use each attribute, as described in Section 3. More specifically, we focus on predicting whether the speaker will use a *spatial relation* to another object (“landmark”). Our motivation for choosing this attribute stems from the fact that previous authors (Viethen and Dale, 2008; Ferreira and Paraboni, 2014) have found substantial variation between different users with respect to their preference towards using spatial relations.

Data We use the GRE3D3 dataset of human-produced REs (Viethen and Dale, 2010), which contains 630 descriptions for 10 scenes collected from 63 users, each describing the same target object in each scene. 35% of the descriptions in this corpus use a spatial relation. An example of such a scene can be seen in Fig. 3.

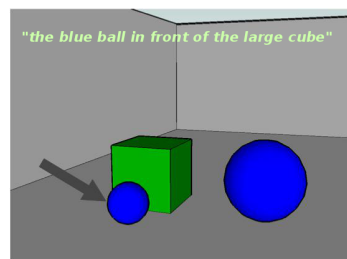


Figure 3: A sample scene with a human-produced RE from the GRE3D3 dataset.

Models We use two baselines for comparison:

Basic: The state-of-the-art model on this task with this dataset, under the assumption that users are seen in training, is presented in Ferreira and Paraboni (2014). They define context features such as type of relation between the target object and its landmark, number of object of the same color or size, etc., then train an SVM classifier to predict the use of each attribute. We recast their model in terms of a log-linear model with the same features, to make it fit with the setup of Section 3.

Basic++: Ferreira and Paraboni (2014) also take speaker features into account. We do not use speaker identity and the speaker’s attribute frequency vector, because we only evaluate on unseen users. We do use their other speaker features (age, gender), together with *Basic*’s context features; this gives us a strong baseline which is aware of manually annotated user group characteristics.

We compare these baselines to our *Group* model for values of K between 1 and 10, using the exact same features as *Basic*. We do not use the speaker features of *Basic++*, because we do not want to rely on manually annotated groups. Note that our results are not directly comparable with those of Ferreira and Paraboni (2014), because of a different training-test split: their model requires having seen speakers in training, while we explicitly want to test our model’s ability to generalize to unseen users.

Experimental setup We evaluate using cross-validation, splitting the folds so that all speakers we see in testing are previously unseen in training. We use 9 folds in order to have folds of the same size (each containing 70 descriptions coming from 7 speakers). At each iteration we train on 8 folds and test on the 9th. At test time, we process each test instance iteratively: we first predict for each instance whether the user u would use a spatial relation or not and test our prediction; we then add the

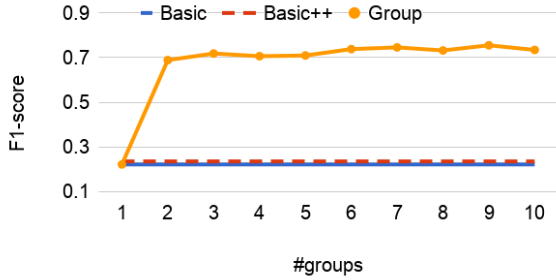


Figure 4: F1 scores on test data for values of K between 1 and 10 in the production experiment.

actual observation from the corpus to the set $D^{(u)}$ of observations for this particular user, in order to update our estimate about their group membership.

Results Figure 4 shows the test F1-score (micro-averaged over all folds) as we increase the number of groups, compared to the baselines. For our *Group* models, these are averaged over all interactions with the user. Our model gets F1-scores between 0.69 and 0.76 for all values of $K > 1$, outperforming both *Basic* (0.22) and *Basic++* (0.23).

In order to take a closer look at our model’s behavior, we also show the accuracy of our model as it observes more instances at test time. We compare the model with $K = 3$ groups against the two baselines. Figure 5 shows that the group model’s F1-score increases dramatically after the first two observations and then stays high throughout the test phase, always outperforming both baselines by at least 0.37 F1-score points after the first observation. The baseline models of course are not expected to improve with time; fluctuations are due to differences between the visual scenes. In the same figure, we plot the evolution of the entropy of the group model’s posterior distribution over the groups (see (7)). As expected, the model is highly uncertain at the beginning of the test phase about which group the user belongs to, then gets more and more certain as the set $D^{(u)}$ of observations from that user grows.

6.2 RE comprehension

Task Our next task is to predict the referent to which a user will resolve an RE in the context of a visual scene. Our model is given a stimulus $s = (r, c)$ consisting of an instruction containing an RE r and a visual context c and outputs a probability distribution over all possible referents b . Such a model can be used by a probabilistic RE generator to select an RE which is highly likely to be correctly understood by the user or predict potential

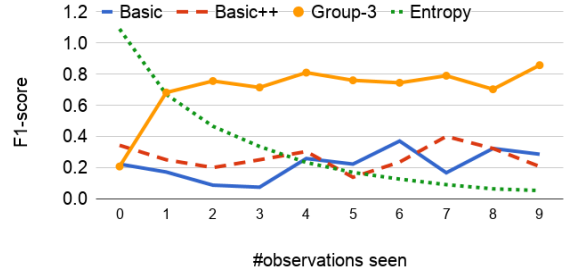


Figure 5: F1-score evolution with increasing number of observations from the user in the production experiment.

misunderstandings (see Section 3).

Data We use the GIVE-2.5 corpus for training and the GIVE-2 corpus for testing our model (the same used by Engonopoulos et al. (2013)). These contain recorded observations of dialog systems giving instructions to users who play a game in a 3D environment. Each instruction contains an RE r , which is recorded in the data together with the visual context c at the time the instruction was given. The object b which the user understood as the referent of the RE is inferred by the immediately subsequent action of the user. In total, we extracted 2927 observations by 403 users from GIVE-25 and 5074 observations by 563 users from GIVE-2.

Experimental setup We follow the training method described in Section 3. At test time, we present the observations from each user in the order they occur in the test data; for each stimulus, we ask our models to predict the referent a which the user understood to be the referent of the RE, and compare with the recorded observation. We subsequently add the recorded observation to the dataset for the user and continue.

Models As a baseline, we use the *Basic* model described in Section 3, with the features of the “semantic” model of Engonopoulos et al. (2013). Those features capture information about the objects in the visual scene (e.g. saliency) and some basic semantic properties of the RE (e.g. color, position). We use those features for our *Group* model as well, and evaluate for K between 1 and 10.

Results on GIVE data *Basic* had a test accuracy of 72.70%, which was almost identical with the accuracy of our best *Group-3* model for $K = 6$ (72.78%). This indicates that our group model does not differentiate between users. Indeed, after training, the 6-group model assigns 81% prior probabil-

ity to one of the groups, and effectively gets stuck with this assignment while testing; the mean entropy of the posterior group distribution only falls from an initial 1.1 to 0.7 after 10 observations.

We speculate that the reason behind this is that the features we use are not sensitive enough to capture the differences between the users in this data. Since our model relies completely on observable behavior, it also relies on the ability of the features to make relevant distinctions between users.

Results on synthetic data In order to test this hypothesis, we made a synthetic dataset based on the GIVE datasets with 1000 instances from 100 users, in the following way: for each user, we randomly selected 10 scenes from GIVE-2, and replaced the target the user selected, so that half of the users always select the target with the highest visual salience, and the other half always select the one with the lowest. Our aim was to test whether our model is capable of identifying groups when they are clearly present in the data and exhibit differences which our features are able to capture.

We evaluated the same models in a 2-fold cross-validation. Figure 6 shows the prediction accuracy for *Basic* and the *Group* models for K from 1 to 10. All models for $K > 1$ clearly outperform the baseline model: the 2-group model gets 62.3% vs 28.6% averaged over all test examples, while adding more than two groups does not further improve the accuracy. We also show in Figure 7 the evolution of the accuracy as $D^{(u)}$ grows: the *Group* model with $K = 2$ reaches a 64% testing accuracy after seeing two observations from the same user. In the same figure, the entropy of the posterior distribution over groups (see production experiment) falls towards zero as $D^{(u)}$ grows. These results show that our model is capable of correctly assigning a user to the group they belong to, once the features are adequate for distinguishing between different user behaviors.

6.3 Discussion

Our model was shown to be successful in discovering groups of users with respect to their behavior, within datasets which present discernible user variation. In particular, if all listeners are influenced in a similar way by e.g. the visual salience of an object, then the group model cannot learn different weights for the visual salience feature; if this happens for all available features, there are effectively no groups for our model to discover.

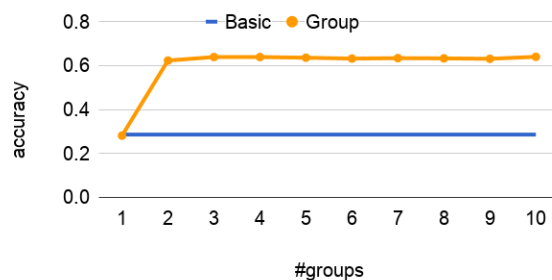


Figure 6: Prediction accuracies in the comprehension experiment with synthetic data.

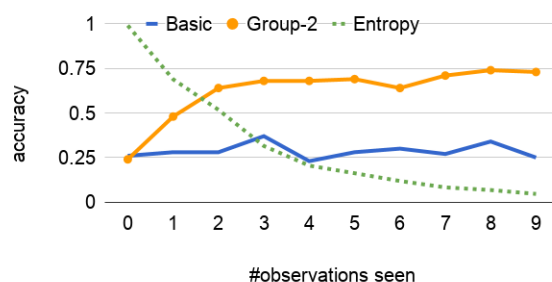


Figure 7: Accuracy evolution with increasing number of observations from the user in the comprehension experiment with synthetic data.

Once the groups have been discovered, our model can then very quickly distinguish between them at test time. This is reflected in the steep performance improvement even after the first user observation in both the real data experiment in 6.1 and the synthetic data experiment in 6.2.

7 Conclusion

We have presented a probabilistic model for NLG which predicts the behavior of individual users of a dialog system by dynamically assigning them to user groups, which were discovered during training². We showed for two separate NLG-related tasks, RE production and RE comprehension, how our model, after being trained with data that is not annotated with user groups, can quickly adapt to unseen users as it gets more observations from them in the course of a dialog and makes increasingly accurate predictions about their behavior.

Although in this work we apply our model to tasks related to NLG, nothing hinges on this choice; it can also be applied to any other dialog-related prediction task where user variation plays a role. In the future, we will also try to apply the basic principles of our user group approach to more sophisticated underlying models, such as neural networks.

²Our code and data is available in <https://bit.ly/2jIu1Vm>

References

- David Benyon and Dianne Murray. 1993. Developing adaptive systems to fit individual aptitudes. In *Proceedings of the 1st international conference on Intelligent user interfaces*. ACM, pages 115–121.
- Adam L. Berger, Stephen A. Della Pietra, and Vincent J. Della Pietra. 1996. A maximum entropy approach to natural language processing. *Computational Linguistics* 22.
- Markus Dräger and Alexander Koller. 2012. Generation of landmark-based navigation instructions from open-source data. In *Proceedings of the Thirteenth Conference of the European Chapter of the ACL*.
- Nikos Engonopoulos and Alexander Koller. 2014. Generating effective referring expressions using charts. In *Proceedings of the INLG and SIGDIAL 2014 Joint Session*. pages 6–15.
- Nikos Engonopoulos, Martín Villalba, Ivan Titov, and Alexander Koller. 2013. Predicting the resolution of referring expressions from user behavior. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. pages 1354–1359.
- Thiago Castro Ferreira and Ivandré Paraboni. 2017. Improving the generation of personalised descriptions. In *Proceedings of the 10th International Conference on Natural Language Generation*. pages 233–237.
- Thiago Castro Ferreira and Ivandr Paraboni. 2014. Referring expression generation: taking speakers preferences into account. In *Proceedings of the International Conference on Text, Speech, and Dialogue*.
- Katja Häuser, Jutta Kray, and Vera Demberg. 2017. Age differences in language comprehension during driving: Recovery from prediction errors is more effortful for older adults. In *Proceedings of CogSci*.
- Srinivasan Janarthanam and Oliver Lemon. 2014. Adaptive generation in dialogue systems using dynamic user modeling. *Computational Linguistics* 40.
- Alexander Koller, Kristina Striegnitz, Andrew Gargett, Donna Byron, Justine Cassell, Robert Dale, Johanna Moore, and Jon Oberlander. 2010. Report on the Second NLG Challenge on Generating Instructions in Virtual Environments (GIVE-2). In *Proceedings of the 6th International Natural Language Generation Conference*.
- Radford M. Neal and Geoffrey E. Hinton. 1999. [A view of the em algorithm that justifies incremental, sparse, and other variants](http://dl.acm.org/citation.cfm?id=308574.308679). In Michael I. Jordan, editor, *Learning in graphical models*, MIT Press, Cambridge, MA, USA, pages 355–368.
- Jorge Nocedal and Stephen Wright. 2006. *Numerical Optimization*. Springer.
- Cecile Paris. 1988. Tailoring object descriptions to a user’s level of expertise. *Computational Linguistics* 14.
- J. A. Rosenblum and J. D. Moore. 1993. Participating in instructional dialogues: Finding and exploiting relevant prior explanations. In *Proceedings of the World Conference on Artificial Intelligence in Education*.
- Kristina Striegnitz, Alexandre Denis, Andrew Gargett, Konstantina Garoufi, Alexander Koller, and Mariet Theune. 2011. Report on the Second Second Challenge on Generating Instructions in Virtual Environments (GIVE-2.5). In *Proceedings of the 13th European Workshop on Natural Language Generation*.
- Jette Viethen and Robert Dale. 2008. The use of spatial relations in referring expression generation. In *Proceedings of the Fifth International Natural Language Generation Conference*. Association for Computational Linguistics, pages 59–67.
- Jette Viethen and Robert Dale. 2010. Speaker-dependent variation in content selection for referring expression generation. In *Proceedings of the Australasian Language Technology Association Workshop 2010*. pages 81–89.
- Marilyn Walker, Amanda Stent, Francois Mairesse, and Rashmi Prasad. 2007. Individual and domain adaptation in sentence planning for dialogue. *Journal of AI Research* 30.