

A Next Utterance Prediction Additional Evaluation Metrics

In Table 5 and Table 6 we show additional results on next utterance prediction. In particular we give results for conditioning on the other person’s persona (“Their”) instead of one’s own (“Self”) or the combination (“Both”). In particular, we see that, knowing the other’s persona does not help precision on this data with these models. Finally, we show in Table 6 for the ranking models the performance difference from training on original personas versus revised ones. Revised personas give improved results, perhaps because the models are forced to learn more than trivial word overlap (i.e. semantic similarity of differing phrases).

B Example Dialogs between Humans and Models

In Tables 7, 8, 9, 11, 12 we show example dialogues between Turkers and various models that were collected as part of the human evaluation of Section 5.2.

C Human Evaluation Measures

After dialogues between humans and a model, we then ask the Turker some additional questions in order to evaluate the quality of the model. They are, in order:

- **Fluency:** We ask them to judge the fluency of the other speaker as a score from 1 to 5, where 1 is “not fluent at all”, 5 is “extremely fluent”, and 3 is “OK”.
- **Engagingness:** We ask them to judge the engagingness of the other speaker *disregarding fluency* from 1-5, where 1 is “not engaging at all”, 5 is “extremely engaging”, and 3 is “OK”.
- **Consistency:** We ask them to judge the consistency of the persona of the other speaker, where we give the example that “I have a dog” followed by “I have no pets” is not consistent. The score is again from 1-5.
- **Profile Detection:** Finally, we display two possible profiles, and ask which is more likely to be the profile of the person the Turker just spoke to. One profile is chosen at random, and the other is the true persona given to the model.

D Profile Prediction

While the main study of this work is the ability to improve next utterance classification by conditioning on a persona, one could naturally consider two tasks: (1) next utterance prediction during dialogue, and (2) profile prediction given dialogue history. In the main paper we show that Task 1 can be improved by using profile information. Task 2, however, can be used to extract such information.

In this section we conduct a preliminary study of the ability to predict the persona of a speaker given a set of dialogue utterances. We consider the dialogues between humans (PERSON 0) and our best performing model, the retrieval-based Key-Value Profile Memory Network (PERSON 1) from Section 5.2. We tested the ability to predict the profile information of the two speakers from the dialogue utterances of each speaker, considering all four combinations. We employ the same IR baseline model used in the main paper to predict profiles: it ranks profile candidates, either at the entire profile level (considering all the sentences that make up the profile as a bag) or at the sentence level (each sentence individually). We consider 100 negative profile candidates for each positive profile, and compute the error rate of predicting the true profile averaged over all dialogues and candidates. The results are given in Table 13, both for the model conditioned on profile information, and the same KV Memory model that is not. The results indicate the following:

- It is possible to predict the humans profile from their dialogue utterances (PERSON 0, Profile 0) with high accuracy at both the profile and sentence level, independent of the model they speaking to.
- Similarly the model’s profile can be predicted with high accuracy from its utterances (PERSON 1, Profile 1) when it is conditioned on the profile, otherwise this is chance level (w/o Profile).
- It is possible to predict the model’s profile from the human’s dialogue, but with a lower accuracy (PERSON 0, Profile 1) as long as the model is conditioned on its own profile. This indicates the human responds to the model’s utterances and pays attention to the model’s interests.

Persona	Method	Original			Revised		
		ppl	hits@1	F1	ppl	hits@1	F1
No Persona		38.08	0.092	0.168	38.08	0.092	0.168
Self Persona	Seq2Seq	40.53	0.084	0.172	40.65	0.082	0.171
	Profile Memory	34.54	0.125	0.172	38.21	0.108	0.170
Their Persona	Seq2Seq	41.48	0.075	0.168	41.95	0.074	0.168
	Profile Memory	36.42	0.105	0.167	37.75	0.103	0.167
Both Personas	Seq2Seq	40.14	0.084	0.169	40.53	0.082	0.166
	Profile Memory	35.27	0.115	0.171	38.48	0.106	0.168

Table 5: **Evaluation of dialog utterance prediction with generative models** in four settings: conditioned on the speakers persona (“self persona”), the dialogue partner’s persona (“their persona”), both or none. The personas are either the original source given to Turkers to condition the dialogue, or the revised personas that do not have word overlap. In the “no persona” setting, the models are equivalent, so we only report once.

Method	No Persona		Self Persona		Their Persona		Both Personas	
	Orig	Rewrite	Orig	Rewrite	Orig	Rewrite	Orig	Rewrite
IR baseline	0.214	0.214	0.410	0.207	0.181	0.181	0.382	0.188
<i>Training on original personas</i>								
Starspace	0.318	0.318	0.481	0.295	0.245	0.235	0.429	0.258
Profile Memory	0.318	0.318	0.473	0.302	0.283	0.267	0.438	0.266
<i>Training on revised personas</i>								
Starspace	0.318	0.318	0.491	0.322	0.271	0.261	0.432	0.288
Profile Memory	0.318	0.318	0.509	0.354	0.299	0.294	0.467	0.331
KV Profile Memory	0.349	0.349	0.511	0.351	0.291	0.289	0.467	0.330

Table 6: **Evaluation of dialog utterance prediction with ranking models** using hits@1 in four settings: conditioned on the speakers persona (“self persona”), the dialogue partner’s persona (“their persona”), both or none. The personas are either the original source given to Turkers to condition the dialogue, or the rewritten personas that do not have word overlap, explaining the poor performance of IR in that case.

- Similarly, the human’s profile can be predicted from the model’s dialogue, but with lower accuracy. Interestingly, the model without profile conditioning is better at this, perhaps because it does not concentrate on talking about itself, and pays more attention to responding to the human’s interests. There appears to be a tradeoff that needs to be explored and understood here.

Overall, the results in this section show that it is plausible to predict profiles given dialogue utterances, which is an important extraction task. Note that better results could likely be achieved with more sophisticated models.

We also study the performance of profile prediction as the dialogue progresses, by computing error rates for dialogue lengths 1 to 8 (the longest length we consider in this work). The results, given in Table 14, show the error rate of predicting the persona decreases in all cases as dialogue length increases.

Persona 1	Persona 2 (Not used by model)
I am a doctor	I eat sardines for breakfast daily
I have one brother	I just had surgery
My mom is a secretary	I love playing Frisbee
I like to eat a lot of meat	I have two large dogs
I am a bodybuilder	I never finished high school

[PERSON 1:] Hi I am Dr. Jenson. What is your name?
[PERSON 2:] Miss Dawes.
[PERSON 1:] Hi Miss Dawes, tell me about your family!
[PERSON 2:] Miss Baggoli, I know it must show on your records that my name is Mary, but everyone has referred to me as Lola since I was a squealing infant in my mother' s arms.
[PERSON 1:] fascinating. My secretary mother would have made note of all of that.
[PERSON 2:] In addition to the remainder, it writes my autobiography.
[PERSON 1:] Oh yeah? well I enjoy a nice steak. how about you?
[PERSON 2:] Father always used to say that if we ever had the money you have... we' d have steak and ice cream three times every day!
[PERSON 1:] I eat all the steak and chicken too, even bacon
[PERSON 2:] Father always used to say that if we ever had the money you have... we' d have steak and ice cream three times every day!
[PERSON 1:] yeah you said that already
[PERSON 2:] Dynamic structure can absorb shock.
[PERSON 1:] yes it can. I know about that as a bodybuilder
[PERSON 2:] Ladies and gentlemen... the strongest man in the world is blowing up a hot water bottle.

Table 7: Example dialog between a human (Person 1) and the OpenSubtitles KV Memory Network model (Person 2).

Persona 1	Persona 2 (Not used by model)
I am afraid of dogs	I am a carpenter
I used to work at a carnival	I used to drink but I stopped five years ago
I like to drink scotch to relax	I have a wife and three kids
I write short stories for a living	I like playing poker
I like dark superhero movies	I dont have many friends

[PERSON 1:] Hi are you relaxing on this fine evening?
[PERSON 2:] I am doing good and you?
[PERSON 1:] I am well, just drinking some scotch and working on my writing
[PERSON 2:] I love to cook and bake but I rarely make time for it
[PERSON 1:] what do you usually bake?
[PERSON 2:] I do not cook, I bake cakes at a local restaurant. how about you?
[PERSON 1:] I used to be a juggler at a carnival but now I am an author
[PERSON 2:] I am a dancer, what do you do?
[PERSON 1:] I just told you. Do you have any pets?
[PERSON 2:] no, I am allergic to pets unfortunately.
[PERSON 1:] I understand. I am not allergic but I am terrified of dogs!
[PERSON 2:] I have a cat. she has been with me through three wives, haha.
[PERSON 1:] Funny that you just told me you have no pets
[PERSON 2:] I do not have any animals.

Table 8: Example dialog between a human (Person 1) and the Seq2Seq model (Person 2).

Persona 1	Persona 2
I am engaged to be married	I am a writer
I used to play sports in high school	I write romance novels
My favorite sport is soccer	I live on a faraway island
I am a college graduate	I am an introvert
I love watching sports on TV	I have two cats

[PERSON 1:] I just got done playing soccer. What have you been up to?
[PERSON 2:] I do not have pets. I spend my free time writing. what about you?
[PERSON 1:] Just graduated from college. Looking for work.
[PERSON 2:] yes I graduated, in writing and poetry
[PERSON 1:] Have you had anything published?
[PERSON 2:] I mainly write fiction novels.
[PERSON 1:] Do you like to watch sports?
[PERSON 2:] do you like kings of leon my favorite by them is use somebody
[PERSON 1:] Are you married? I will be married soon.
[PERSON 2:] haha, no time. I have got a novel to finish.
[PERSON 1:] What is your novel about?
[PERSON 2:] it is a love story. I really enjoy writing. do you write or read?
[PERSON 1:] Like to read. Spend most time watching football.
[PERSON 2:] yes I do I am and author for a living

Table 9: Example dialog between a human (Person 1) and the Key-Value Profile Memory Network with Self Persona.

Persona 1	Persona 2
I dont like exercise or physical activity	My favorite type of food is Italian
My favorite food is ice cream	I have a big pet dog
I am a homebody	I want to travel to Europe
I have a husband who I love very much	I love reading books about travel
I love watching movies and TV	I like to walk long distances

[PERSON 1:] Hello, how are you doing today?
[PERSON 2:] I am good. I just got finished walking my dog.
[PERSON 1:] Oh cool! I am just sitting watching a movie.
[PERSON 2:] I have been watching that one before. I love it.
[PERSON 1:] Me too. I don't like going out much so I enjoy movies and TV.
[PERSON 2:] that is a great idea. I am going to read a book on writing.
[PERSON 1:] Oh, are you a writer?
[PERSON 2:] yes, I love to read. I am a great author.
[PERSON 1:] What sorts of topics do you write about?
[PERSON 2:] I could read all of those. I am sure you enjoy reading.
[PERSON 1:] I do, but I prefer movies.
[PERSON 2:] I have not, I can not do that. I have to go to the east coast.
[PERSON 1:] Why can't you watch movies?
[PERSON 2:] I have not. I have always had that done.

Table 10: Example dialog between a human (Person 1) and the Generative Profile Memory Network with Self Persona.

Persona 1	Persona 2 (Not used by model)
I am very active	I am an engineering student
I live in the desert	I work at Eureka
I enjoy the outdoors	I like to celebrate with friends
I was named something pretty common	I enjoy spending time with my toes in the sand
I instruct people how to get fit as a job	I study to be a businessman

[PERSON 1:] Hello there, how are you?
[PERSON 2:] I'm fine.
[PERSON 1:] How is the weather, it is hot here
[PERSON 2:] It's so hot.
[PERSON 1:] oh yes tell me about it especially when you live in a desert.
[PERSON 2:] the sea is shining.
[PERSON 1:] you at the beach, i love the beach
[PERSON 2:] wait, wait, wait...
[PERSON 1:] what's up, you going somewhere
[PERSON 2:] you can be a good friend.
[PERSON 1:] yes, if you want me to
[PERSON 2:] I'm sorry.

Table 11: Example dialog between a human (Person 1) and the Language Model trained on the Open-Subtitles 2018 dataset (does not use persona).

Persona 1	Persona 2 (Not used by model)
I love ordering from Amazon	I really enjoy watching classic films over the new ones
I am a mother	My mother is the greatest chef in the world
I have a very active imagination	I stand out among my surroundings
Winter holidays have always been my favorite	My son and daughter lived with me full time after I
I love the Harry Potter books	separated from my husband

[PERSON 1:] Good afternoon! How are you?
[PERSON 2:] i'm not sure. I was just talking about the other day.
[PERSON 1:] Do you like to read? I like reading Harry Potter.
[PERSON 2:] I love you so much
[PERSON 1:] That's funny. You don't know me.
[PERSON 2:] I think you have to be a good person.
[PERSON 1:] Where do you live? Are you in the US?
[PERSON 2:] yes, I just found out that there were only 10 people who were in the city.
[PERSON 1:] Do you have any children? I am a mother to 1 cat.
[PERSON 2:] this is the first time in history, but not a few.
[PERSON 1:] Is it cold where you are?
[PERSON 2:] I don't even know what I'm talking about.

Table 12: Example dialog between a human (Person 1) and the Language Model trained on the Twitter dataset (does not use persona).

Speaker	Profile	Profile Level		Sentence Level	
		KV Profile	KV w/o Profile	KV Profile	KV w/o Profile
PERSON 0	Profile 0	0.057	0.017	0.173	0.141
PERSON 0	Profile 1	0.234	0.491	0.431	0.518
PERSON 1	Profile 0	0.254	0.112	0.431	0.349
PERSON 1	Profile 1	0.011	0.512	0.246	0.530

Table 13: **Profile Prediction.** Error rates are given for predicting either the persona of speaker 0 (Profile 0) or of speaker 1 (Profile 1) given the dialogue utterances of speaker 0 (PERSON 0) or speaker 1 (PERSON 1). This is shown for dialogues between humans (PERSON 0) and either the KV Profile Memory model (“KV Profile”) which conditions on its own profile, or the KV Memory model (“KV w/o Profile”) which does not.

Speaker	Profile	Dialogue Length							
		1	2	3	4	5	6	7	8
PERSON 0	Profile 0	0.76	0.47	0.35	0.29	0.23	0.19	0.17	0.17
PERSON 0	Profile 1	0.51	0.39	0.32	0.29	0.27	0.27	0.25	0.25
PERSON 1	Profile 0	0.57	0.52	0.48	0.46	0.45	0.43	0.43	0.43
PERSON 1	Profile 1	0.81	0.58	0.48	0.47	0.45	0.44	0.43	0.43

Table 14: **Profile Prediction By Dialog Length.** Error rates are given for predicting either the persona of speaker 0 (Profile 0) or of speaker 1 (Profile 1) given the dialogue utterances of speaker 0 (PERSON 0) or speaker 1 (PERSON 1). This is shown for dialogues between humans (PERSON 0) and the KV Profile Memory model averaged over the first N dialogue utterances from 100 conversations (where N is the “Dialogue Length”). The results show the accuracy of predicting the persona improves in all cases as dialogue length increases.