Natural Language Engineering **6** (3–4): 363–377. Printed in the United Kingdom

© 2000 Cambridge University Press

Towards developing general models of usability with PARADISE

MARILYN WALKER, CANDACE KAMM
and DIANE LITMAN
AT&T Labs – Research, 180 Park Avenue, Florham Park, NJ 07932-0971, USA
e-mail: {walker,cak,diane}@research.att.com

(*Received 3 November 1999*)

Abstract

The design of methods for performance evaluation is a major open research issue in the area of spoken language dialogue systems. This paper presents the PARADISE methodology for developing predictive models of spoken dialogue performance, and shows how to evaluate the predictive power and generalizability of such models. To illustrate the methodology, we develop a number of models for predicting system usability (as measured by user satisfaction), based on the application of PARADISE to experimental data from three different spoken dialogue systems. We then measure the extent to which the models generalize across different systems, different experimental conditions, and different user populations, by testing models trained on a subset of the corpus against a test set of dialogues. The results show that the models generalize well across the three systems, and are thus a first approximation towards a general performance model of system usability.

1 Introduction

The development of methods for evaluating, comparing and predicting system performance is a major open research issue in the area of Spoken Language Dialogue Systems (SLDSs). In 1997, we proposed PARADISE (PARAdigm for DIalogue System Evaluation) as a general integrative framework for evaluating SLDSs (Walker, Litman, Kamm and Abella 1997b). PARADISE addressed three research goals:

- to support the comparison of multiple systems or multiple versions of the same system doing the same domain tasks;
- to provide a method for developing predictive models of the usability of a system as a function of a range of system properties; and
- to provide a technique for making generalizations across systems about which properties of a system impact usability, i.e. to figure out 'what really matters to users'.

In previous work, we conducted a number of within-system comparisons, based on the application of PARADISE to experimental dialogues collected with three different systems: ANNIE, an SLDS for voice dialing and messaging (Kamm,

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

Litman and Walker 1998); ELVIS, an SLDS for accessing email (Walker, Fromer and Narayanan 1998b; Walker, Fromer, Di Fabbrizio, Mestel and Hindle 1998a); and TOOT, an SLDS for accessing online train schedules (Litman, Pan and Walker 1998; Litman and Pan 1999). We also used the models we learned for the ELVIS system to provide an objective performance function for learning optimal dialogue strategies using reinforcement learning (Walker *et al.* 1998b; Walker 2000).

This paper focuses on the predictive power and generalizability of the models learned by applying PARADISE. We first develop a number of models for predicting system usability (as measured by user satisfaction). Then we measure the extent to which these models generalize across different systems, different experimental conditions, and different user populations, by testing models trained on a subset of the corpus against a test set of dialogues. The experimental corpus consists of 544 ANNIE, ELVIS and TOOT dialogues, totaling almost 42 hours of speech.

The structure of the paper is as follows. We first review the PARADISE evaluation framework in Section 2. Then Section 3 describes the ANNIE, ELVIS and TOOT systems and the experimental setup for collecting the corpus of dialogues. Section 4 presents results from developing predictive models of user satisfaction and testing these models on subsets of the corpus. We then discuss the results and suggest areas for future work.

2 PARADISE

The state of the art in evaluating SLDSs, prior to the proposal of the PARADISE framework, was to evaluate the SLDS in terms of a battery of both subjective and objective metrics. Some of these metrics focused on task completion or transaction success. Others were based on the performance of the SLDS's component technologies, such as speech recognizer performance (Sparck Jones and Galliers 1996; Hirschman, Dahl, McKay, Norton and Linebarger 1990; Ralston, Pisoni and Mullennix 1995; Pallett 1985). Subjective metrics included measures of user satisfaction (Shriberg, Wade and Price 1992), or ratings generated by dialogue experts as to how cooperative the system's utterances were (Bernsen, Dybkjaer and Dybkjaer 1996).

The motivation for PARADISE was that problems arise when a battery of metrics is used. First, several different metrics may contradict one another. For example, Danieli and Gerbino compared two train timetable systems (Danieli and Gerbino 1995), and found that one version of their SLDS had a higher transaction success rate and produced fewer inappropriate and repair utterances, but that the other version produced dialogues that were approximately half as long. However, they could not report whether the higher transaction success or the length of the dialogue was more critical to performance, or what tradeoffs between them might be acceptable.

Secondly, to make generalizations across different systems performing different tasks, it is important to know how multiple factors impact performance and how users' perceptions of system performance depend on the dialogue strategy, and on tradeoffs among other factors like efficiency, usability and accuracy (Sparck-Jones and Galliers 1996).

The PARADISE framework derives a combined performance metric for a dia-

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

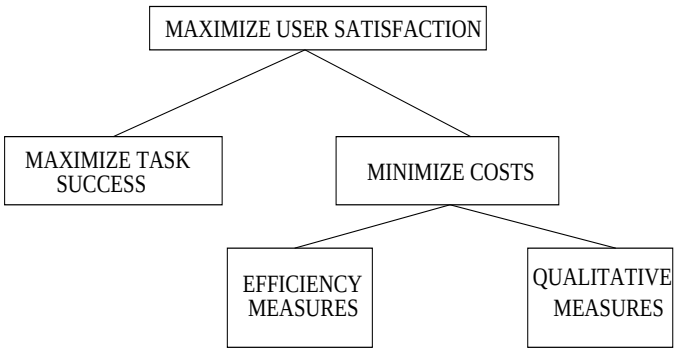


Fig. 1. PARADISE's structure of objectives for spoken dialogue performance.

logue system as a weighted linear combination of a task-based success measure and dialogue costs. To specify what factors should go into this combined performance metric, PARADISE posits a particular model of performance, illustrated in Figure 1. The model proposes that the system's primary objective is to maximize user satisfaction. Task success and various costs that can be associated with the interaction are both contributors to user satisfaction. Dialogue costs are of two types: dialogue efficiency and quality. Efficiency costs are measures of the system's efficiency in helping the user complete the task, such as the number of utterances to completion of the dialogue. Dialogue quality costs are intended to capture other aspects of the system that may have strong effects on user's perception of the system, such as the number of times the user had to repeat an utterance to make the system understand.[1]

Applying PARADISE to dialogue data requires that dialogue corpora be collected via controlled experiments during which users subjectively rate their satisfaction. In addition, the other components of the model, i.e. task success and interaction costs, must be either automatically logged by the system or be hand-labeled. The PARADISE model posits that a performance function can then be derived by applying multivariate linear regression with user satisfaction as the dependent variable and task success, dialogue quality and dialogue efficiency measures as independent variables.

Modeling user satisfaction as a function of task success and dialogue cost metrics is intended to lead to predictive performance models for SLDSs, so that values for user satisfaction could be predicted on the basis of a number of simpler metrics that can be directly measured from the system logs, without the need for extensive experiments with users to assess user satisfaction. To make this predictive use of PARADISE a reality, the models that are derived from experiments with one set of systems or user populations should be *generalizable* to other systems or other user populations. This paper reports results from applying PARADISE to three different systems, to begin to show what generalizations can be made across systems and user populations.

[1] Developing qualitative cost measures that can be measured automatically is an active area of research (Hirschman and Pao 1993; Walker, Hirschman and Aberdeen 2000).

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

3 Experimental design

3.1 Experimental setup

All of the experiments applying PARADISE used a similar experimental setup. In each experiment, human subjects carried out a dialogue using one of the three dialogue systems. Each system was implemented using a general-purpose platform for phone-based spoken dialogue systems (Kamm, Narayanan, Dutton and Ritenour 1997). The platform consisted of a speech recognizer that supports barge-in so that the user can interrupt the system when it is speaking. It also provided an audio server for both voice recordings and Text-To-Speech (TTS), an interface between the computer running the system and the telephone network, a module for application specific functions, and modules for specifying the application grammars and the dialogue manager.

The dialogues were obtained in controlled experiments. The ANNIE experiments were designed to evaluate (1) the effect of novice versus expert user populations, and (2) the impact of a short tutorial on the novice user population (Kamm *et al.* 1998). The ELVIS experiments were designed (1) to evaluate the effect of different dialogue strategies for managing initiative and presentation of information (Walker *et al.* 1997a, 1998a); and (2) to conduct experiments applying reinforcement learning to dialogue so that the ELVIS system could learn from experience which dialogue strategies are optimal (Walker *et al.* 1998b; Walker 2000). The TOOT1 experiments were designed to evaluate the effects of (1) task difficulty and (2) cooperative versus literal response strategies (Litman *et al.* 1998). The TOOT2 experiments studied the effect of user-adapted interaction in conjunction with different initiative and confirmation dialogue strategies (Litman and Pan 1999).

All of the experiments required users to complete a set of application tasks in conversations with a particular version of the system. Following PARADISE, these tasks were characterized in terms of particular items of information that the user must find out from the system. Each set of tasks was also designed to be representative of typical tasks in that domain.

Experiments with ANNIE were designed to be comparable with ELVIS so the users in the ANNIE experiments performed the same tasks as those for ELVIS. Each user had to perform three tasks in sequence in three difference conversations with the system. To examine the effect of a short tutorial, some users also performed an additional task as part of a tutorial interaction before starting the experimental tasks. The example email access task below was used for both ANNIE and ELVIS:

- You are working at home in the morning and plan to go directly to a meeting when you go into work. Kim said she would send you a message telling you where and when the meeting is. Find out **the Meeting Time** and **the Meeting Place**.

To complete this task with ANNIE, users accessed their email by either linearly navigating through an email folder, as in the sample dialogue in Figure 2, or by using referring expressions like *message 2* to jump back to messages that had been heard already. ANNIE's prompts were pre-recorded, and ANNIE would switch

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

from pre-recorded prompts to text-to-speech when reading the message headers and the bodies of the messages.

A sample dialogue illustrating the completion of the same task with the system-initiative version of ELVIS is shown in the sample dialogue in Figure 3. In ELVIS, users could ask to hear a summary of their inbox, or ask to go ahead and read their messages. Once reading messages was initiated, users were prompted for message selection criteria, which allowed them to select messages based on their content by specifying values from the sender or the subject field. Since users could select a message directly rather than needing to linearly navigate through the folder, it was possible for users of ELVIS to complete the task more efficiently. All of ELVIS's utterances were realized via text-to-speech.

A. Annie here.
U. Get my messages.
A. First message. "This is a message from Michael about Evaluation Technical memo"...
(*user barges in*)
U. Next message.
A: Message 2. "This is a message from Kim about interviewing Antonio.
I'd like to interview him. I could also go along to lunch. Kim."
2:16 P.M. Monday, March
23rd. What's next?
U: Next message.
A: Message 3. "This is a message from Kim about meeting today. The
meeting this morning is at 10:00 in 2d 516. See you there, Kim."
(*user barges in*)
U: I'm done here.

Fig. 2. Sample dialogue with the ANNIE system.

For TOOT, each user performed four tasks in sequence. Each task was represented by a scenario where the user had to find a train satisfying certain constraints, by using the system to retrieve and process online train schedules. A sample dialogue with the TOOT system for the following task scenario is in Figure 4. This dialogue illustrates the use of a mixed-initiative and implicit confirmation dialogue strategy.

- Try to find a train going to **New York City** from **Philadelphia** on Sunday at 10:30 pm. If you cannot find an exact match, find the one with the **closest** departure time. Please write down the **exact departure time** of the train you found as well as the **total travel time**.

The experiments resulted in 108 dialogues with the ANNIE system, 268 dialogues with the ELVIS system, and 168 dialogues with the TOOT system for a total of 544 dialogues, consisting of almost 42 hours of speech.

3.2 Data collection

This section discusses in detail the metrics that were logged and hand-labeled in experiments with all three systems. We have not previously made specific recommendations about what metrics to use, as this is an active area of research (Walker,

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

E: Hi, ELVIS here. You have 5 new and 0 unread messages in your inbox. Say Read, Summarize, or say Help for more options.
U: Summarize.
E: In your top level inbox, from Kim, there's 1 message about "Lunch." and 1 message about "Meeting Today". From Michael, there's 1 message about "Evaluation group meeting." From Noah, there's 1 message about "Call Me Tomorrow" and 1 message about "Interviewing Antonio." And from Owen, there's 1 message about "Agent Personality."
U: Read.
E: Select by Sender, Subject or in order of arrival?
U: Sender.
E: Which Sender?
U: Kim.
E: In the messages from Kim, there's 1 message about "Interviewing Antonio" and 1 message about "Meeting Today." The first message is titled, "Interviewing Antonio." (*user barges in*)
U: Next message.
A: "This is a message from Kim about meeting today. The meeting this morning is at 10:00 in 2d 516. See you there, Kim."
U: I'm done here.

Fig. 3. Sample dialogue with ELVIS.

T: Hi this is AT&T Amtrak Schedule System. This is TOOT. How may I help you?
U: I'd like to get a train from Philadelphia to New York.
T: I heard you say go to New York from Philadalphia. When do you want to leave?
U: Sunday.
T: I heard you say leave on Sunday. What time do you want to leave?
U: 10:30 pm.
T: I heard you say leave around 10:30 pm. Do you want me to find the trains from Philadelphia to New York on Sunday around 10:30 pm now?
U: Yes.
T: I am going to get the train schedule for you...

Fig. 4. Sample dialogue with the TOOT system.

Hirschman and Aberdeen 2000; Polifroni, Seneff, Glass and Hazen 1998; Hirschmann and Pao 1993; Danieli and Gerbino 1995; Bernsen *et al.* 1996). However, any metric for task success, dialogue efficiency or dialogue quality previously proposed in the literature can be easily incorporated into the PARADISE framework.

The metrics discussed below combine dialogue quality and efficiency measures, with a focus on measures that can be automatically logged or computed. Three different methods were used to collect the data: (1) All of the dialogues were recorded; (2) The dialogue manager logged the system's dialogue behavior and a number of other measures discussed below; (3) Users filled out web page forms after each task (task success and user satisfaction measures). Measures are summarized in Figure 5 and described in more detail below.

The **dialogue efficiency** metrics were calculated from the dialogue recordings and the system logs. The length of the recording was used to calculate the elapsed time

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

- **Dialogue Efficiency Metrics**
  - — elapsed time, system turns, user turns
- **Dialogue Quality Metrics**
  - — mean recognition score, timeouts, rejections, helps, cancels, bargeins (raw)
  - — timeout%, rejection%, help%, cancel%, bargein% (normalized)
- **Task Success Metrics**
  - — task completion as per survey
- **User Satisfaction**
  - — the sum of TTS Performance, ASR Performance, Task Ease, Interaction Pace, User Expertise, System Response, Expected Behavior, Comparable Interface, Future Use.

Fig. 5. Metrics collected for spoken dialogues.

in seconds (**ET**) from the beginning to the end of the interaction. Measures for the number of **System Turns**, and the number of **User Turns**, were calculated on the basis of the system logging everything it said and everything it heard the user say.

The **dialogue quality** measures were derived from the recordings, the system logs and hand-labeling. A number of system behaviors that affect the quality of the resulting dialogue were automatically logged. These included the number of timeout prompts (**timeouts**) played when the user didn't respond as quickly as expected, the number of recognizer rejections (**rejects**) where the system's confidence in its understanding was low and it said something like *I'm sorry I didn't understand you.* User behaviors that the system perceived that might affect the dialogue quality were also logged: these included the number of times the system played one of its context specific help messages because it believed that the user had said *Help* (**helps**), and the number of times the system reset the context and returned to an earlier state because it believed that the user had said *Cancel* (**cancels**). The recordings were used to check whether users barged in on system utterances, and these were labeled on a per-state basis (**bargeins**).

Another measure of dialogue quality was recognizer performance over the whole dialogue, calculated in terms of concept accuracy. The recording of the user's utterance was compared with the logged recognition result to calculate a concept accuracy measure for each utterance by hand. Concept accuracy is a measure of semantic understanding by the system, rather than word for word understanding. For example, the utterance *Read my messages from Kim* contains two concepts, the *read* function, and the *sender:kim* selection criterion. If the system understood only that the user said *Read*, then concept accuracy would be 0.5. Mean concept accuracy was then calculated over the whole dialogue and used, in conjunction with ASR rejections, to compute a Mean Recognition Score **MRS** for the dialogue.

Because the goal is to generate models that will generalize across systems and tasks, we also thought it important to introduce metrics that are likely to generalize. All of the efficiency metrics seemed unlikely to generalize since, e.g. the **elapsed time** to complete of a task depends on how difficult the task is. Other research suggested that the dialogue quality metrics were more likely to generalize (Litman

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

- Was the system easy to understand in this conversation? (**TTS Performance**)
- In this conversation, did the system understand what you said? (**ASR Performance**)
- In this conversation, was it easy to find the message you wanted? (**Task Ease**)
- Was the pace of interaction with the system appropriate in this conversation? (**Interaction Pace**)
- In this conversation, did you know what you could say at each point of the dialogue? (**User Expertise**)
- How often was the system sluggish and slow to reply to you in this conversation? (**System Response**)
- Did the system work the way you expected him to in this conversation? (**Expected Behavior**)
- In this conversation, how did the system's voice interface compare to the touch-tone interface to voice mail? (**Comparable Interface**)
- From your current experience with using the system to get your email, do you think you'd use the system regularly to access your mail when you are away from your desk? (**Future Use**)

Fig. 6. User satisfaction survey for ELVIS and ANNIE experiments.

et al. 1999), but it seemed likely that the raw counts would be task specific. Thus we normalized the dialogue quality metrics by dividing the raw counts by the total number of utterances in the dialogue. This resulted in the **timeout%, rejection%, help%, cancel%** and **bargein%** metrics.

The web page forms are the basis for calculating Task Success and User Satisfaction measures. Users reported their perceptions as to whether they had completed the task (**Comp**).[2] They also had to provide objective evidence that they had in fact completed the task by filling in a form with the information that they had acquired from the system.[3]

To calculate User Satisfaction, users were asked to evaluate the system's performance with a user satisfaction survey. The survey probed a number of different aspects of the users' perceptions of their interaction with the SLDS in order to focus the user on the task of rating the system, as in Shriberg *et al.* (1992), Jack, Foster and Stentiford (1992), Love, Dutton, Foster, Jack and Stentiford (1994) and Larsen (1999). A sample survey is in Figure 6. The surveys used for the three SLDSs were identical except that the **Comparable Interface** question was eliminated from the TOOT survey. The surveys were multiple choice and each survey response was mapped into the range of 1 to 5. Then the values for all the responses were summed, resulting in a **User Satisfaction** measure for each dialogue ranging from 8 to 40.

[2] *Yes, No* responses are converted to *1,0*.

[3] This supports an alternative way of calculating **Task Success** objectively by using the Kappa statistic to compare the information that the users filled in with a key for the task (Walker *et al.* 1997b). However earlier results indicated that user's *perception* of task success was a better predictor of overall satisfaction, so here we simply use perceived task success as measured by **Comp**.

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

4 Results of applying PARADISE

Table 1 summarizes the experimental corpus of 544 dialogues in terms of experimental conditions and the metrics that were collected. This section reports results for

- training models to predict user satisfaction via multivariate linear regression, and
- testing these models across different systems, different experimental conditions, and different user populations to determine to what extent they generalize.

Table 1. *Performance measure means per dialogue for different SLDSs*

Measure	SI[a]	MI[b]	Exps[c]	NovTut[d]	Novs[e]	Fixed[f]	Adapt[g]
Comp	0.87	0.80	1.0	0.96	0.73	0.75	0.87
User Turns	21.5	17.0	10.8	13.5	21.4	13.6	15.3
Sys Turns	24.2	21.2	11.7	15.4	25.5	14.0	15.7
ET[h]	339.14 s	296.18 s	156.5 s	195.6	280.3 s	238.6 s	234.1 s
MRS[i]	.88	.72	.80	.74	.67	.77	.78
TimeOuts	2.65	4.15	.64	1.28	2.4	0.32	0.25
TimeOut%	0.11	0.19	0.04	0.08	0.10	0.02	0.02
Cancs	0.34	0.02	0.17	0.17	0.61	1.02	0.42
Canc%	0.01	0	0.01	0.01	0.03	0.07	0.02
Helps	0.67	0.92	0.08	0.75	3.3	0.26	0.05
Help%	0.03	0.05	0.01	0.05	0.14	0.02	0.00
BargeIns	3.6	3.6	4.5	4.5	5.8	1.0	1.6
BargeIn%	0.08	0.09	0.2	0.17	0.15	0.08	0.10
Rejects	0.86	1.58	1.8	3.3	6.6	1.1	1.1
Reject%	0.04	0.08	0.17	0.22	0.25	0.07	0.08
USAT[j]	28.9	25.0	32.8	30.8	23.0	28.3	31.0

[a] ELVIS system initiative.
[b] ELVIS mixed initiative.
[c] ANNIE experts.
[d] ANNIE novices who had a tutorial.
[e] ANNIE novices.
[f] TOOT1 and non-adaptable TOOT2.
[g] Adaptable TOOT2.
[h] Elapsed time.
[i] Mean recognition score.
[j] User satisfaction.

An overall summary of the results for testing and training predictive models is in Table 2. Each model discussed was trained by performing a stepwise multivariate linear regression on some set of dialogues with user satisfaction as the dependent variable and the other metrics discussed above and shown in Table 1 as the independent variables. For each model we describe the training set, indicate the performance of the trained model on the training set, describe the test set, and indicate the performance of the trained model on the test set. Table 3 lists the factors that were significant predictors of user satisfaction for each of these models.

We first trained models for each system on a random 90% of the dialogues for that system. Results for each system are given in the first three rows of Table 2, showing

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

Table 2. *Testing the predictive power of different models*

Training set[a]	R^2 Training (SE)[b]	Test set[c]	R^2 Test (SE)[d]
ANNIE 90%	0.50 (0.009)	ANNIE 10%	0.40 (0.07)
ELVIS 90%	0.39 (0.003)	ELVIS 10%	0.43 (0.03)
TOOT 90%	0.56 (0.014)	TOOT 10%	0.54 (0.05)
ALL 90%	0.47 (0.004)	ALL 10%	0.50 (0.04)
ELVIS 90%	0.42	TOOT	0.55
ELVIS 90%	0.42	ANNIE	0.36
NOVICES	0.47	ANNIE EXPERTS	0.04
MRS $\leq$ 0.95	0.46	MRS $>$ 0.95	0.23

[a] The training set used to train the model.
[b] The amount of variance in user satisfaction in the training set accounted for by the trained model. When 10 fold cross-validation is used to estimate R^2, the value reported is a mean and the standard error (SE) is shown.
[c] A characterization of the test set.
[d] The amount of variance in user satisfaction in the test set accounted for by the trained model.

the degree to which the model is able to predict user satisfaction in the training set (Column labeled R^2 Training).[4] We then tested whether a model trained on a random 90% of the corpus for a particular system could predict user satisfaction in the remaining 10% of the dialogues for that system. The results in the first three rows of Table 2 suggest that all the models generalize well to the test data, although the model fits for ELVIS and ANNIE on the training data are worse than that for TOOT.[5]

We then examined models trained on the combined data for all three systems. The ALL row in Table 2 gives the results for training on a random 90% of the dialogues from the combined corpus and testing on the remaining 10% of the dialogues. As the table shows, the model is a good predictor of user satisfaction in the test set, accounting for 50% of the variance in R^2.

Next we examined whether models trained on a random 90% of one of the systems (ELVIS) could predict user satisfaction for users using one of the other systems (ANNIE, TOOT). The results for this test are shown in rows 5 and 6 of Table 2. Surprisingly, the ELVIS model accounts for 55% of the variance in user satisfaction for the TOOT1 and TOOT2 data sets, accounting for more of the variance in TOOT than it does on the ELVIS test set, and fitting the TOOT test data as well as a model trained on TOOT data (Row 3). This suggests that it is easier to find a good model for the TOOT corpus than for the other two corpora.

[4] For the first four rows of the table, 10-fold cross-validation is used to estimate R^2. Thus, the training and testing R^2 values represent means, and standard errors (SE) are shown. For the rest of the table, the R^2 values are obtained using a single training set and a separate held-out test set.
[5] However, as we discuss in more detail below, the predictive capability of a model is dependent on the size of the training set and there are fewer ANNIE dialogues in the corpus.

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

Table 3. *Factors that have the most predictive power in different models*

Training set[a]	Factors[b]
ANNIE 90%	help%, mrs, comp, bargein%
ELVIS 90%	mrs, comp, reject%, bargein%
TOOT 90%	mrs, comp, et, reject%
ALL 90%	mrs,comp,reject%, help%, bargein%, et
NOVICES	mrs, comp, reject%, bargein%, help%, et, timeout%
MRS $\leq$ 0.95	mrs, comp, reject%, help%, et, bargein%, timeout%

[a] The training set used to train the model.
[b] Significant predictors in order of contribution.

The ELVIS model also accounts for 36% of the variance in user satisfaction in the ANNIE data set (Row 6).

We then turn to the question of whether models trained on one user population generalize to a different type of user population. We examined this in two ways. First, we had a group of expert users who participated in the ANNIE experiments. These users had been using the ANNIE system for over 6 months on a regular basis. All of the other users who participated in the experiments were novice users who were not familiar with the limitations of spoken dialogue systems. Out of the 544 dialogues, 36 were dialogues with expert users and the remaining 508 dialogues were collected with novice users. We trained a model with the dialogues from the novice users and tested it against the dialogues with the experts. The results are shown in Table 2 in the NOVICES row. The results show that models trained on novice users *do not* easily generalize to an expert user population. The correlation of the predicted values to the actual values is only 0.19; the novice model only accounts for 4% of the variance in user satisfaction in the expert population.

Secondly, speech recognition performance varies a great deal from user to user, and thus it is important to assess whether the models trained on ANNIE, ELVIS and TOOT would generalize to systems whose speech recognition performance overall was much better. In order to examine this question, we selected the 93 out of 544 of the dialogues that had a Mean Recognition Score (MRS) better than 0.95 (GoodASR dialogues). We used the 451 dialogues where MRS was less than or equal to 0.95 to train a model for predicting user satisfaction and then tested it on the GoodASR dialogues. The results shown in the MRS $\leq$ 0.95 row of Table 2 show that the model only accounts for 23% of the variance in R^2 in the test data. This is not surprising since the contribution of MRS to user satisfaction is large in the model trained on the rest of the corpus.

Table 3 also shows which factors were found to be significant predictors of user satisfaction in each of the trained models (Factors column), ordered by degree of contribution. Except for the ANNIE system model, recognizer performance (MRS) and task success (COMP) are always the largest contributors to user satisfaction, followed by the percentage of the dialogue turns that were rejections (REJECT%), where the system asked the user to repeat or paraphrase what they had just said. The

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

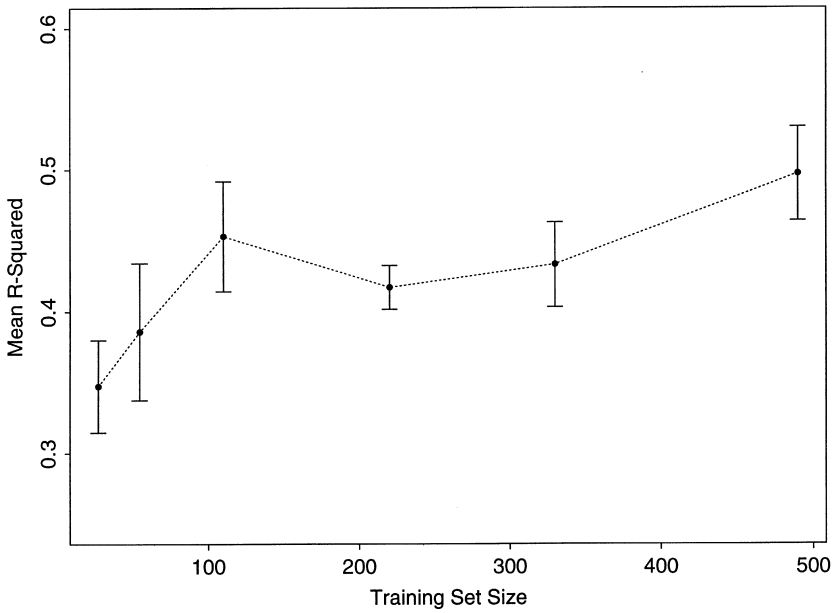


Fig. 7. Model fit as a function of training set size.

models for the ANNIE experiments reflect the fact that the population of users was deliberately varied in these experiments. ANNIE experts had been using ANNIE for a period of about 6 months and many of them had considerable experience with other spoken dialogue systems. The ANNIE model reported in row 1 shows that the percentage of requests for *help* was a significant predictor of user satisfaction, reflecting the fact that some users needed help significantly more than others, and these same users typically had much lower satisfaction with the system.

Finally, we examined how training set size affects performance by applying cross-validation to the combined data from all three systems. We trained with 10 random samples of 27, 55, 110, 220, 330 and 490 dialogues and then tested on an unseen 10% of the dialogues. Figure 7 plots model fit (R^2) as a function of the amount of data used to train the model. The figure shows that the model fit is starting to level out as the training set size approaches 200 dialogues. While the mean for training on 490 dialogues is greater than the mean for training on 110 dialogues, the error bars for the model fits overlap, suggesting that the use of larger corpora may not increase the the predictive power of the models.

5 Discussion and future work

When we introduced the PARADISE framework for evaluating Spoken Language Dialogue Systems (SLDSs) we hoped to apply an empirical approach to developing general models of system usability. We proposed that user satisfaction could be used as a measure of system usability, and that it would be possible to develop models for predicting usability by learning which of range of other metrics were good predictors of user satisfaction.

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

This paper focused on developing general models for predicting user satisfaction based on experimental data collected with three different SLDSs. The results showed that: (1) a model developed on 90% of the dialogues for ANNIE accounted for 50% ($\pm$ 1%) of the variance in the training data and 40% ($\pm$ 7%) of the variance in the test data; (2) a model developed on 90% of the dialogues for ELVIS accounted for 39% ($\pm$ 0.3%) of the variance in the training data and 43% ($\pm$ 3%) of the variance in the test data; and (3) a model developed on 90% of the dialogues for TOOT accounted for 56% ($\pm$ 1%) of the variance in the training data and 54%($\pm$ 5%) of the variance in the test data. In other words, the performance of each model on the unseen 10% of the dialogues was comparable to performance on the training set.

We then showed that a general model, trained on 90% of the dialogues combined from all three systems accounted for 47% ($\pm$ 0.4%) of the variance in user satisfaction, and that this model accounted for 50% ($\pm$ 4%) of the variance in user satisfaction in the unseen 10% of the dialogues. In other words, the performance of the general model trained on 90% of the data is statistically indistinguishable on the test set and the training set.

We also tested the extent to which a model trained on one system (ELVIS) can predict user satisfaction for the other two systems (ANNIE and TOOT). These tests showed that the ELVIS model could account for 55% of the variance in user satisfaction in TOOT and 36% of the variance in user satisfaction in ANNIE. The ELVIS model actually does better at predicting user satisfaction in TOOT than it does on a random unseen 10% of ELVIS dialogues, and its predictive power on the ANNIE data is no worse than a model trained on 90% of the ANNIE dialogues. This suggests that there are general factors that predict user satisfaction that are shared by the models developed on each of the three systems; this finding is also supported by the fact that the models share many predictive factors as shown in Table 3.

We also examined the extent to which dialogues trained on one user population generalize to dialogues collected with other user populations. We first tested models trained on dialogues with novices against dialogues with experts. The novice models did not predict user satisfaction in the expert population. This shows that it is important in future work to carry out longitudinal studies of users as they gain experience with the system, and that general models must be trained on a sample of dialogues including both expert and novice users. Finally, the results were also weaker when we tested models trained on dialogues with users who experienced average recognizer performance against dialogues with users who experienced excellent recognizer performance.

The results reported contribute to the goal of developing general models of system usability, however there are several ways in which this work should be extended. First, it would be useful to repeat this analysis with a large corpus of cross-system dialogues collected via participation in the Darpa Communicator evaluations (Walker *et al.* 2000). Secondly, it is important to extend these techniques to dialogues collected in field studies, as opposed to the controlled experimental environment described here, perhaps by first examining the differences between dialogues collected with predefined scenarios for a particular user to dialogues collected when the user is

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

asked to define their own task, as in Baggia, Castagneri and Danieli (1998). Finally, it is important to determine whether it is possible to achieve better model fits than those reported here. The model fits range from the ELVIS model that accounts for 39% of the variance in user satisfaction to the TOOT model that accounts for 56% of the variance in user satisfaction. One possibility is that larger dialogue corpora might help, but Figure 7 suggests that the use of larger dialogue corpora may not, in and of itself, automatically improve the models. However, this result may be misleading since there is a fair amount of homogeneity in the dialogue systems examined here. While they varied tasks, interaction style, and dialogue strategies, all three systems were built with the same platform and applied similar approaches to dialogue manager specification.

There are also at least two other possible avenues for improving the models of user satisfaction. One line of attack is investigating whether modeling techniques that are more powerful than linear models, could improve predictive performance. While we show (Walker *et al.* 2000) that the model fits for tree models (Brieman, Friedman, Olshen and Stone 1984) are not significantly better than linear models, other modeling techniques could be applied. Secondly, it may be possible to develop a number of additional metrics for predicting user satisfaction that would generalize across systems (Walker *et al.* 2000). The results presented here suggest that users are highly attuned to the quality of the interaction rather than efficiency, but it is a challenge to develop metrics that assess the quality of the interaction that can be reliably hand-labelled or automatically measured. For example, the field is only just beginning to develop metrics that assess the quality of the system's strategies for presenting information and for providing context-sensitive help. In future work, we hope to contribute to the development of automatic metrics that are highly correlated with the quality of the user's experience when interacting with a spoken dialogue system.

References

Baggia, P., Castagneri, G. and Danieli, M. (1998) Field trials of the italian arise train timetable system. *Interactive Voice Technology for Telecommunications Applications, IVTTA*, pp. 97–102.

Bernsen, N. O., Dybkjaer, H. and Dybkjaer, L. (1996) Principles for the design of cooperative spoken human-machine dialogue. *Int. Conference on Spoken Language Processing (ICSLP 96)*, pp. 729–732.

Brieman, L., Friedman, J. H., Olshen, R. A. and Stone, C. J. (1984) *Classification and Regression Trees.* Wadsworth and Brooks, Monterey, CA.

Danieli, M. and Gerbino, E. (1995) Metrics for evaluating dialogue strategies in a spoken language system. *Proc. AAAI Spring Symposium on Empirical Methods in Discourse Interpretation and Generation*, pp. 34–39.

Hirschman, L., Dahl, D. A., McKay, D. P., Norton, L. M. and Linebarger, M. C. (1990) Beyond class A: A proposal for automatic evaluation of discourse. *Proc. Speech and Natural Language Workshop*, pp. 109–113.

Hirschman, L. and Pao, C. (1993) The cost of errors in a spoken language system. *Proc. Third European Conference on Speech Communication and Technology*, pp. 1419–1422.

Jack, M., Foster, J. C. and Stentiford, F. W. (1992) Intelligent dialogues in automated telephone services. *Int. Conference on Spoken Language Processing (ICSLP)*, pp. 715–718.

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press

Kamm, C., Litman, D. and Walker, M. A. (1998) From novice to expert: The effect of tutorials on user expertise with spoken dialogue systems. *Proc. International Conference on Spoken Language Processing (ICSLP98)*.

Kamm, C., Narayanan, S., Dutton, D. and Ritenour, R. (1997) Evaluating spoken dialog systems for telecommunication services. *5th European Conference on Speech Technology and Communication, EUROSPEECH 97*, pp. 2203–2206.

Larsen, L. B. (1999) Combining objective and subjective data in evaluation of spoken dialogues. *ESCA Workshop on Interactive Dialogue in Multi-Modal Systems*, pp. 89–92.

Litman, D. J. and Pan, S. (1999) Empirically Evaluating an Adaptable Spoken Dialogue System. *Proc. 7th International Conference on User Modeling*.

Litman, D. J., Pan, S. and Walker, M. A. (1998) Evaluating Response Strategies in a Web-Based Spoken Dialogue Agent. *Proc. ACL/COLING 98: 36th Annual Meeting of the Association of Computational Linguistics*, pp. 780–787.

Litman, D. J., Walker, M. A. and Kearns, M. J. (1999) Automatic detection of poor speech recognition at the dialogue level. *Proc. 37th Annual Meeting of the Association of Computational Linguistics (ACL99)*, pp. 309–316.

Love, S., Dutton, R. T., Foster, J. C., Jack, M. A. and Stentiford, F. W. M. (1994) Identifying salient usability attributes for automated telephone services. *Int. Conference on Spoken Language Processing (ICSLP)*, pp. 1307–1310.

Pallett, D. S. (1985) Performance assessment of automatic speech recognizers. *J. Res. Nat. Bureau of Standards*, **90**: 371–387.

Polifroni, J., Seneff, S., Glass, J. and Hazen, T. (1998) Evaluation methodology for a telephone-based conversational system. *Proc. 1st International Conference on Language Resources and Evaluation*, Granada, Spain, pp. 42–50.

Ralston, J. V., Pisoni, D. B. and Mullennix, J. W. (1995) Perception and comprehension of speech. In: A. K. Syrdal, R. W. Bennet and S. L. Greenspan, editors, *Applied Speech Technology*, pp. 233–287. CRC Press.

Shriberg, E., Wade, E. and Price, P. (1992) Human-machine problem solving using spoken language systems (SLS): Factors affecting performance and user satisfaction. *Proc. DARPA Speech and NL Workshop*, pp. 49–54.

Sparck-Jones, K. and Galliers, J. R. (1996) *Evaluating Natural Language Processing Systems*. Springer-Verlag.

Walker, M., Fromer, J., Di Fabbrizio, G., Mestel, C. and Hindle, D. (1998a) What can I say?: Evaluating a spoken language interface to email. *Proc. Conference on Computer Human Interaction (CHI 98)*, pp. 582–589.

Walker, M., Hindle, D., Fromer, J., Di Fabbrizio, G. and Mestel, C. (1997a) Evaluating competing agent strategies for a voice email agent. *Proc. European Conference on Speech Communication and Technology, EUROSPEECH97*.

Walker, M., Kamm, C. and Boland, J. (2000) Developing and testing general models of spoken dialogue system performance. *Proc. Language Resources and Evaluation Conference, LREC-2000*.

Walker, M. A. (2000) An application of reinforcement learning to dialogue strategy selection in a spoken dialogue system for email. *J. Artificial Intelligence Research*, **12**: 387–416.

Walker, M. A., Fromer, J. C. and Narayanan, S. (1998b) Learning optimal dialogue strategies: A case study of a spoken dialogue agent for email. *Proc. 36th Annual Meeting of the Association of Computational Linguistics, COLING/ACL 98*, pp. 1345–1352.

Walker, M. A., Hirschman, L. and Aberdeen, J. (2000) Evaluation for DARPA communicator spoken dialogue systems. In *Proc. 2nd International Conference on Language Resources and Evaluation*, Athens, Greece.

Walker, M. A., Litman, D., Kamm, C. A. and Abella, A. (1997b) PARADISE: A general framework for evaluating spoken dialogue agents. *Proc. 35th Annual Meeting of the Association of Computational Linguistics, ACL/EACL 97*, pp. 271–280.

https://doi.org/10.1017/S1351324900002503 Published online by Cambridge University Press