A Method for Transferring Probabilistic User Models between Environments

David L. Roberts[1] and Fred Roberts[2]

[1] Department of Computer Science, North Carolina State University
robertsd@csc.ncsu.edu
[2] DIMACS, Rutgers University
froberts@dimacs.rutgers.edu

Abstract. Chief among the inputs to decision making algorithms in narrative or game environments is a model of player or opponent decision making. A challenge that will always face designers is to specify that model ahead of time, when actual data from the environment is likely not to be available. Absent corpora of data, designers must intuit these models as best they can, incorporating domain or expert knowledge when available. To make this process more precise, we derive a theoretically grounded technique to transfer an observed user model from one domain to another. We answer the question: "How can a model obtained from observations of one environment inform a model for another environment?" We verify the accuracy of our techniques using data from previous user studies.

1 Introduction

Interactive narrative experiences are marked by two important characteristics: 1) a large space of player interactions, some subset of which are specified as aesthetic goals by authors; and 2) the affordance for players to express self-agency and interact in a meaningful way. As a result, players are (often unknowing) participants in the creation of their narrative experience. They cannot be assumed to be cooperative, nor to be adversarial. Thus, researchers have designed computational paradigms that work with players to co-create experiences without the need for goal-oriented models of player behaviors. To effectively work with a (possibly unknowing) partner, systems rely on various types of models that describe player behaviors. In this paper, we will look at probabilistic models. More specifically, we will rigorously examine how *what we learn in one narrative or game environment can be put to use in a new environment (sometimes even a new game), thereby alleviating the requirement for authors to create a new player model from scratch*. We will describe a way in which we can use quantitative or probabilistic data from one domain to estimate the change in probability in a second domain that a given alternative is preferred to another after an action has been applied to the first. Further, this will allow us to use data from other domains to develop player models for interactive narratives without the need for detailed hand authoring of the models.

The earliest work on interactive narrative is traceable to Laurel [7], who first proposed the idea of a human Drama Manager (DM) to adaptively guide actors on stage to bring about a better narrative experience for audience members. A human director can

M. Si et al. (Eds.): ICIDS 2011, LNCS 7069, pp. 43–54, 2011.

adapt their (possibly tacit) knowledge about how audiences may react to narrative adaptions in a variety of settings. Humans are very good at transferring knowledge between domains, *e.g.*, by learning a mental model of someone else's behavior in one setting and then applying that knowledge in another (possibly unrelated) setting.

After Laurel's initial work, progress in drama management shifted almost exclusively to algorithmic developments for computational realizations of DMs in virtual environments (see Mateas [10] and Roberts & Isbell [17] for a surveys). Unfortunately, while humans may be good at knowledge transfer between domains, computers are generally poor at transferring knowledge between domains. Therefore, the job of specifying decision making rules for different domains falls upon humans, making the lessons a computer can learn in one domain useless in other domains.

Notable among early efforts to move virtual drama management beyond "script-and-trigger" decision making are those of Bates [1] and Weyhrauch [23]. They published work on what ultimately became known as Declarative Optimization-based Drama Management (DODM) [11]. A DODM instance relies upon a number of author-specified components of the narrative environment and decision making processes, most notably (for our purposes) a probabilistic model of player behavior and a set of actions the DM can take to effect changes in player decisions. This model describes how players are likely to transition through a narrative environment, given the history of the story and DM actions. In practice, this can be very difficult to implement effectively [19,20,23].

In this paper, we derive a mathematical psychology-inspired method for transferring the observed effectiveness of a DM action from one narrative domain to another. To our knowledge, this is the first application of these mathematical psychology models to game design. It will enable authors to design their models using data from previous experiences, rather than having to engineer them based on intuition or educated guesses.

While our interest in drama management in general and the DODM formalism more specifically led us to investigate this problem, our method is not limited to narrative domains. In fact, any setting where a probabilistic model of player behavior serves as input into a decision making algorithm can benefit from increased specificity. Despite this generality, our approach is not universally applicable—there is no silver bullet. Below, we will discuss conditions that might let our method yield accurate results.

2 DODM and Related Work

First proposed by Weyhrauch [23] as a problem of pseudo-adversarial search, DODM has been studied in recent years as a formal decision making process [11,12,15,18]. A DODM drama manager is characterized by four components: 1) a set of **plot events with precedence constraints**; 2) a set of **DM actions** that *operate on plot events*; 3) a probabilistic **player model** of the progression of events that encodes the likely behavior of players; and 4) an **evaluation function** encoding author specified goals.

The player model is typically represented as a probability function $P(S'|A,S)$ which represents the probability of a story event S' occurring immediately subsequent to a story event S given that the DM has executed action A after story event S has occurred. Actions are typically modeled as "operating on plot events" and generally take one of three forms: *deny* an event, *cause* an event, or *hint* at an event. From our perspective, deniers and causes aren't particularly interesting as their outcome is deterministic; however, the result of a hint, depending on how that hint is delivered, may change significantly. Hence the need for a detailed model of probabilistic action effects. If $P(S'|S)$ is the probability that event S' is chosen by a player after event S, and A operates on S', then likely $P(S'|A, S) \neq P(S'|S)$. How much different is a question traditionally answered by authors' estimates. The main contributions of this paper are two models that enable this probability change to be estimated in a principled, data-driven way.

To date, the vast majority of the work on DODM has been abstract, so actions have represented more general methods for adjusting the likelihood a player will experience a story event they operate on. For example, in the original formulation due to Weyhrauch [23], the transition model was hand-authored based on a "best guess" of how effective a class of actions might be. This best guess used a uniform distribution over successive story events as a base case. When actions were applied to a particular story event, the weight on that event was increased and the distribution recalculated by normalizing the weights.

More recently, Sullivan *et al.* have examined two other types of player models that encapsulate "world knowledge" about the story environment [19,20]. Based on the Manhattan distance (or L_1 norm), these models attribute *a priori* weights to plot events based on the physical distance between the plot events and the player in the story environment. In order to extract probabilities, weights are normalized (as in the uniform approach). When a DM action operates on a plot event, the weight can be updated according to changes in the Manhattan distance between the player and the event trigger.

Lastly, Roberts *et al.* [14,16] have investigated the use of social psychology influence techniques [2] as a framework for actions, delegating the implementation to a separate process. By defining actions as the application of social influence in the environment, data from existing studies in other domains that estimate probabilistic changes in players' behaviors can be used to make similar estimates in a new domain. An advantage over earlier approaches is that there is a vast literature describing evaluations of the effectiveness of influence methods in various real-life settings that can be leveraged to implement a player model. In addition to data collected from other virtual domains, data from the psychology literature can be input into the algorithm we present below. We are asking: "If we learn from one domain that an action has a certain measurable effect on the probability that a player will prefer one alternative to another, what does that tell us about the effect the action will have on the probability of preferring the first alternative to a third in another domain?" For example, imagine players in an MMORPG are faced with the choice between two quests and that they prefer the first with probability 0.4. Further, suppose we know that when the "default effect" [6] is applied to the first quest as a hint they choose it with probability 0.7. Lastly, suppose we know players tend to prefer a third quest over a fourth with probability 0.5. We now want to know the probability they will prefer the first quest over a third if the default effect is applied to it. Our method will rigorously answer that question.

3 A Method for Cross-Domain Transfer of Probabilities

Here we describe our technique for transferring probabilistic data between domains.

To make things concrete, suppose $A, B, C, \ldots$ represent potential outcomes and we use a subscript 0 as in A_0 to denote the outcome in the base situation and 1 as in A_1 to denote the outcome when an action has been applied to it. In the context of an interactive narrative, these outcomes would be story events and the action applied would be a hint. Recalling the above example of quests in an MMORPG, the outcomes would be alternative quests and the action would be applying the default effect to the choice. We will use *alternative* to mean an outcome with a treatment and we will use letters X and Y when we deal with alternatives and don't specify whether they involve a base condition or condition with an applied action. We will use the letters A, B, C to represent outcomes before we specify the condition. We would like to go from two outcomes for which we have preference data in the base case to the probability the first outcome with an action applied to it is preferred to the second. More generally, we would like to calculate the probability $P(A_1 > C_0)$, interpreted as the probability that outcome A under the influence of an action is preferred to outcome C if no action is applied, if we know $P(A_0 > C_0)$, *i.e.*, the probability that outcome A is preferred to outcome C both without actions applied. This would give us a way to estimate the effect of applying an action to an outcome.

We will discuss how to obtain such probabilities based on the types of data available in the literature or that we might obtain from our own data collection. To use the model we present here, we have to assume that the model of utility and preference we base our work on accurately describes how people choose between alternatives and that the magnitude of the effect of an action strategy observed on one set of alternatives will be similar when applied to other alternatives. We acknowledge that our assumptions may not always hold; however, even if they don't hold our approach can provide a starting point that authors can tweak if their intuition tells them the assumptions are violated.

3.1 Types of Data Available

Over the years, there have been countless social psychology studies published that describe the effects of influence on real-life situations. The results of those studies are generally reported as either *numerical data* or *probabilistic data*. Both of these data are easily obtained from experiments run on narrative or game environments as well.

The numerical data experiments report findings based on quantities or scales that are directly measured. For example, Folkes *et al.* present an experiment measuring the effects of product scarcity on usage [4]. Specifically, they conducted experiments where participants were given a measured amount of shampoo in various sized containers. The results of the experiment indicate that under certain conditions the more perceived scarcity, the more the product is used by the study participants (*e.g.*, 500 ml of shampoo in a 1,000 ml bottle leads to 87 ml of use on average whereas 250 ml of shampoo in a 1,000 ml bottle leads to 121 ml of use [4]).

Another example of numerical data is that of Regan [13], who examined how reciprocity in the form of a favor can lead to increased levels of compliance. Regan's measurements were of the quantity of lottery tickets purchased under different conditions. When he reported that in the base condition study participants bought on average 1.00 lottery tickets whereas participants in the influence condition bought on average 1.91 lottery tickets, Regan was able to show reciprocity's significant effect on quantity.

On the other hand, data are sometimes reported using frequencies or probabilities. In that case, the effects of the different study conditions induce a probability distribution over outcomes, or give us a way to obtain the probability that one outcome is preferred to another. This type of data is a more natural fit with the DODM probabilistic transition model and, in fact, one of our approaches will be to translate count into frequency data. Consider the effects of reciprocity discussed by Cialdini [2]. He reports that the Disabled American Veterans Association gets a response rate of approximately 18% when soliciting donations via a mass mailing campaign. When reciprocity is invoked via an unsolicited gift being included in the mass mailing, the donation response rate rises to 35%. This is an example of probabilistic data reported in the literature.

Another example of probabilistic data is reported by Cialdini *et al.* in their study of the effect of reciprocal concessions on compliance with requests for volunteers. In that case, it was found that a mere 16.7% of study participants agreed to volunteer in the control condition, but when reciprocal concessions were employed 50.0% agreed.

3.2 The Strict Utility Model for Numerical Data

When data is given in terms of quantities rather than frequencies or probabilities, some models allow us to compute frequencies or probabilities. For example, in 1929, Zermelo proposed what has come to be known as the strict utility model [24]. This model, widely studied in the mathematical psychology literature, describes probabilistic choice in a forced-choice pair comparison system, where for every pair of alternatives X and Y, each trial asks a subject to decide if they prefer X to Y or Y to X, with no indifference allowed. Then $P(X > Y)$ represents the frequency with which (the probability that) X is preferred to Y. We say that a pair comparison system satisfies the ***strict utility model*** if and only if there is a utility function over the alternatives f that satisfies:

$$P(X > Y) = \frac{f(X)}{f(X) + f(Y)}. \tag{1}$$

This model will form the basis for one method through which we transfer numerical results from one domain to probabilities of player choice in another domain. The strict utility model is not applicable in every situation, and in Section 4 we present experimental results to verify the model. Let us consider the example from Folkes *et al.* [4] discussed above. In those results, there are five different conditions for which data are presented; however, here we will focus on two of them: A_0 and A_1, which according to our notation represent a control (A_0) in which no action (or in this case influence) is applied and treatment condition (A_1) in which an influence action is used. Suppose that the data reported for each outcome is given by a function q such that $q(A_0)$ is the quantity reported for A_0, $q(A_1)$ the quantity reported for A_1, etc. In addition to assuming the strict utility model, we assume the utility of an alternative X is proportional to the quantity reported for that alternative: $f(X) = \lambda \cdot q(X)$. While it need not be the case that λ is equal for all X in every scenario, our experimental results indicate that this assumption leads to reasonably accurate results in many cases.

We are interested in what $q(A_0)$ and $q(A_1)$ tell us about a player's probabilistic choice in a different domain. Specifically, we will show that based on the quantity data

$q(A_0), q(A_1)$, if we know $P(A_0 > C_0)$ for some C, then we can derive $P(A_1 > C_0)$. Suppose that $P(A_0 > C_0) = p$. According to the strict utility model we have:

$$p = P(A_0 > C_0) = \frac{f(A_0)}{f(A_0) + f(C_0)} \quad (2)$$

$$\implies f(C_0) = \frac{f(A_0)}{p} - f(A_0) \quad (3)$$

Thus, we have:

$$P(A_1 > C_0) = \frac{f(A_1)}{f(A_1) + f(C_0)} \quad (4)$$

$$= \frac{f(A_1)}{f(A_1) + \frac{f(A_0)}{p} - f(A_0)} \quad (5)$$

$$= \frac{q(A_1)}{q(A_1) + \frac{q(A_0)}{p} - q(A_0)} \quad (6)$$

This therefore allows us to calculate $P(A_1 > C_0)$ given that we know $q(A_0)$ and $q(A_1)$ from previously collected data and we either know from the literature or assume the value of $P(A_0 > C_0)$. Thus, we can construct a probabilistic transition model based on the assumption of probabilities without actions.

To see how this method is applied, we use the above example from Regan [13]. We have $q(A_0) = 1.00$ and $q(A_1) = 1.91$ when reciprocity is applied. Suppose C is another outcome for which we have no reason to think it or A would be preferred *a priori*. Thus, we assume $P(A_0 > C_0) = 0.5$ in the new domain. We have:

$$P(A_0 > C_0) = p = 0.5$$

$$P(A_1 > C_0) = \frac{1.91}{1.91 + \frac{1.00}{0.5} - 1.00} = 0.656,$$

an estimate of the effect of reciprocity on two different alternatives.

3.3 The Fechnerian Utility Model for Probabilistic Data

In certain cases, the available data are presented as frequencies or probabilities rather than quantities and therefore cannot be interpreted as (proportional to) utility estimates. In such cases, the strict utility model (Equation 1) does not apply. Instead, we turn to a more general model of forced-choice pair comparisons known as the ***Fechnerian utility model*** [3,9]. The Fechnerian utility model, which like the strict utility model is also widely studied in the mathematical psychology literature, holds if there is a monotone increasing function $\phi : \mathbb{R} \to \mathbb{R}$ so that for all outcomes X, Y,

$$P(X > Y) = \phi[f(X) - f(Y)] \quad (7)$$

As before, $f(X)$ represents the utility of X. Therefore $P(X > Y)$ is a function of the difference between the utilities of X and Y. As with the strict utility model, while it need not be the case that such a function ϕ exists in every scenario, our experimental results indicate that this assumption leads to reasonable results in many cases.

Often times it assumed that ϕ is a cumulative distribution function and $P(X > Y)$ is then interpreted as the probability that X has higher utility than Y. Thurstone's early work on the topic [21,22] made the assumption that ϕ followed a standard normal (Gaussian distribution with $\mu = 0, \sigma^2 = 1$):

$$P(X > Y) = \phi[f(X) - f(Y)] = \int_{-\infty}^{[f(X)-f(Y)]} N(x)dx$$

Later, Guilford [5] and Luce [8] proposed the logistic distribution as a better model:

$$P(X > Y) = \phi[f(X) - f(Y)] = \frac{1}{1 + e^{-[f(X)-f(Y)]}} \tag{8}$$

Note that

$$\phi(x) = \frac{1}{1 + e^{-x}} = z \tag{9}$$

$$\implies x = -ln\left(\frac{1}{z} - 1\right) \tag{10}$$

Assuming the Guilford-Luce special case of the Fechnerian utility model, we derive a method for transferring probabilistic data from one domain to another. For this model and a given action, we have two known values, one unknown value, and a value supplied by authors. For two alternatives A and B, we know from existing literature or another source the probabilities that A is preferred to B in the source domain both with and without an action applied to A (specified by $P(A_0 > B_0)$ and $P(A_1 > B_0)$ respectively). The author supplies as input to the model the base probability $P(A_0 > C_0) = p$ indicating the preference of A over C they expect to see in the prediction domain. Using these three values, we can compute $P(A_1 > C_0)$ which is an estimate of the effectiveness of the action in the prediction domain.

Using the Fechnerian utility model with the logistic distribution, we have:

$$P(A_0 > B_0) = \phi[f(A_0) - f(B_0)] = \frac{1}{1 + e^{-[f(A_0)-f(B_0)]}} \tag{11}$$

$$P(A_1 > B_0) = \phi[f(A_1) - f(B_0)] = \frac{1}{1 + e^{-[f(A_1)-f(B_0)]}} \tag{12}$$

Let $\alpha = f(A_0) - f(B_0)$ and $\gamma = f(A_1) - f(B_0)$. Note that we know these values from Equations 10, 11, and 12. Suppose we know the probability A_0 is preferred to C_0 is equal to p, for some $p \in (0, 1]$. That is $P(A_0 > C_0) = p$. Further, let $\beta = f(A_0) - f(C_0) - \alpha$, so $\alpha + \beta = f(A_0) - f(C_0)$. Note that we know β since we know α and since the Guilford-Luce version of the Fechnerian utility model gives us $\alpha + \beta$.

Now we have

$$\begin{aligned} P(A_1 > C_0) &= \phi[f(A_1) - f(C_0)] \\ &= \phi[f(A_1) - f(B_0) + f(B_0) - f(C_0)] \\ &= \phi[\gamma + \beta] \end{aligned}$$

Since we know the value of γ and β, we know the effect of the action in the prediction domain, *i.e.*, $P(A_1 > C_0)$. Note that the above method works for the general Fechnerian utility model. It is just not as easy to get x from $\phi(x)$.

Using the above data from Cialdini [2], we have $P(A_0 > B_0) = 0.18$ which implies $\alpha = f(A_0) - f(B_0) = -1.51635$. Additionally, we have $P(A_1 > B_0) = 0.35$, which implies that $\gamma = f(A_1) - f(B_0) = -0.61904$. Let us find an outcome C so that in the control situation, we have no reason to prefer either A or C, *i.e.*, so that $P(A_0 > C_0) = p = 0.5$. Then we have $f(A_0) - f(C_0) = 0$ and, further, $\beta = f(A_0) - f(C_0) - \alpha = 1.51635$. Therefore,

$$P(A_1 > C_0) = \phi[\gamma + \beta] = \frac{1}{1 + e^{-[0.89371]}} = 0.71040$$

which makes intuitive sense.

4 A Sample of Numerical Results

To evaluate the effectiveness of our approach, we have run our two models using data from a number of sources. There are three parameters for our evaluation:

1. The *source data domain* provides probabilities or counts that serve as input into the strict or Fechnerian utility model, specifically $q(A_0)$ and $q(A_1)$ (strict utility model) or $P(A_0 > B_0)$ and $P(A_1 > B_0)$ (Fechnerian utility model); we will use three sources, one coming from a published paper, "*paper*", (details below) and the other two from published user studies, "*study 1*", and "*study 2*" (details below).
2. The *baseline estimate domain* provides the estimate of the preference over alternatives when no action applies, *i.e.*, $P(A_0 > C_0)$; we will use one of three sources of data, a pure "*guess"* and observations from "*study 1*" or "*study 2*".
3. The *prediction domain* is where we will gather data to determine how accurate our model is based on input data; we will compare outcomes, *i.e.*, predicted vs. observed probability $P(A_1 > C_0)$, with data from *study 1* and *study 2*.

Data for the source domain from "paper" comes from Folkes, Martin, and Gupta [4] for count data and Cialdini [2] for probabilistic data. Respectively, the $P(A_0 > B_0)$ and $P(A_1 > B_0)$ values for these are: 87, 121 for count data and 0.18, 0.35 for frequency data. The data for "study 1" and "study 2" come from published interactive storytelling systems [14,16]. Respectively the $P(A_0 > B_0)$ and $P(A_1 > B_0)$ values for these are: 52, 81 and 71, 85 for count data and 0.515, 0.808 and 0.707, 0.851 for probabilistic data. Unless otherwise specified, the "guess" input was $P(A_0 > C_0) = 0.5$.

The data from Study 1 and Study 2 were collected from a web-based choose your own adventure storytelling system. There were a number of differences between the domains, including the story itself (the setting, characters, *etc.*). Both systems utilized a branching narrative with forced choice two-alternative decisions that players were presented with. Actions in both domains were natural language utterances that were added to the story text, and designed to invoke the social psychological principle of scarcity [2]. In both stories, scarcity was the "strategy" for the action, but there were multiple concrete realizations. Specifically, there were four unique scarcity utterances in Study 1 and three unique scarcity utterances in Study 2. Thus, the input data above and analysis below is based on the average of the applications of all scarcity actions (27 times total in Study 1 and 74 times total in Study 2).

Table 1. A comparison of the predicted and observed probabilities $P(A_1 > C_0)$ using the *strict utility model* with count data under various input conditions. The "parameters" column indicates whether a paper (P), guess (G), study 1 (S_1), or study 2 (S_2) was used for the source domain, baseline estimate, or prediction domain respectively.

parameters	predicted	observed	error
P, G, S_1	0.5817	0.8077	0.2250
P, G, S_2	0.5817	0.8514	0.2696
P, S_1, S_1	0.5964	0.8077	0.2113
P, S_2, S_2	0.7700	0.8514	0.0813
P, S_1, S_2	0.5964	0.8514	0.2549
P, S_2, S_1	0.7700	0.8077	0.0377
S_1, S_1, S_2	0.6234	0.8514	0.2280
S_1, S_2, S_2	0.7895	0.8514	0.0619
S_2, S_2, S_1	0.7424	0.8077	0.0653
S_2, S_1, S_1	0.5609	0.8077	0.2478

First, consider the data in Table 1 where the results of our strict utility model for count data are presented under various conditions. In that Table, the predicted and observed probabilities $P(A_1 > C_0)$ are compared and their error (absolute difference between the observed and expected probabilities) is listed as well. We found that results were varied. In cases where S_2 was used as our baseline for $P(A_0 > C_0)$ average error was very low (0.0615); however, when either a guess or S_1 was the baseline, average error was notably higher (0.2478 and 0.2355 respectively). These contributed largely to the overall average error (and standard deviation) we observed of 0.1684 (0.0939).

Of particular interest in Table 1 are the rows where the prediction domain (third parameter) is not equal to either the source domain or baseline input (parameters one and two) because these tests are indicative of a transfer between two completely different domains. In two of those cases, *i.e.*, P, S_2, S_1 and S_2, S_2, S_1, the results are very good.

Next, consider the data in Table 2 where the results of our Fechnerian utility model are presented under various conditions. First, note that the error rate is notably lower using this version of our model with a mean (and standard deviation) of 0.0735 (0.0426) compared to 0.1684 (0.0939) for the strict utility model. Second, notice that similarly to the strict utility model, in general the best performance occurred when S_2 was used as the baseline again. In other words, our observations suggest that a large influence on the overall accuracy of our models is the baseline guess for $P(A_0 > C_0)$. As before, the rows of particular interest are those where the prediction domain is distinct from the source domain and input baseline. In most of those cases the results are very good.

To examine this effect more closely, consider the data presented in Figure 1. Those data were obtained by holding fixed the source and prediction domains and varying the baseline guess $P(A_0 > C_0)$ from 0.1 to 0.9 in increments of 0.1. We used S_1 as the source domain to predict performance in S_2 and *vice versa*. There are a few interesting

Table 2. A comparison of the predicted and observed probabilities $P(A_1 > C_0)$ using the *Fechnerian utility model* with probability data under various input conditions. The "parameters" column indicates whether a paper (P), guess (G), study 1 (S_1), or study 2 (S_2) was used for the source domain, baseline estimate, or prediction domain respectively.

parameters	predicted	observed	error
P, G, S_1	0.7104	0.8077	0.0973
P, G, S_2	0.7104	0.8514	0.1410
P, S_1, S_1	0.7227	0.8077	0.0850
P, S_2, S_2	0.8552	0.8514	0.0038
P, S_1, S_2	0.7227	0.8514	0.1286
P, S_2, S_1	0.8552	0.8077	0.0475
S_1, S_1, S_2	0.8077	0.8514	0.0437
S_1, S_2, S_2	0.9050	0.8514	0.0536
S_2, S_2, S_1	0.8514	0.8077	0.0437
S_2, S_1, S_1	0.7165	0.8077	0.0912

things to note about this figure. First, the true observed value of $P(A_0 > C_0)$ is 0.707 for the S_2 prediction domain and 0.515 for the S_1 prediction domain. In both sets of data reported in Figure 1, the Fechnerian utility model produced lowest error estimate with a baseline of 0.6, in between the 0.515 [16] and 0.707 [14] true values. Note, because the strict utility model uses count data, we can't make the same comparison. Also note that the Fechnerian utility model was generally more accurate than the strict utility model with a lower error in seven of the nine test cases.

Lastly, and perhaps most significantly, notice that change between the S_1, G, S_2 and S_2, G, S_1 series within a model is minimal. This strongly suggests the biggest source of variance in the accuracy of our model is not the transfer between domains, but the required input from the author about the baseline, *i.e.*, $P(A_0 > C_0)$. In other words, as long as the data from the source domain is an accurate representation of the effect of the action, the results in the prediction domain will be accurate provided the author's baseline is reasonably accurate. This is extremely encouraging for the use of this model. While we have not eliminated the need for authors to provide accurate information, we have reduced the amount of information through the use of our models, requiring only data from a source domain and a baseline guess on the player behavior.

5 Extension to n-Choice Alternatives

Although we have only described the strict and Fechnerian utility models for two-choice alternative situations, they can easily be applied in n-choice alternatives as well. The basic idea is to "leave one out" and consider the remaining alternatives as one "composite" alternative. For example, suppose there are five alternatives $A^1, \ldots, A^5$ and we are interested in knowing the effect of applying an action to A^3. We can use as input into our model $P(A_0^3 > \{A_0^1, A_0^2, A_0^4, A_0^5\})$ and $P(A_1^3 > \{A_0^1, A_0^2, A_0^4, A_0^5\})$. Here, if $\mathbf{A} = \{A^j\}$ we will define $f(\mathbf{A}) = \sum_{A^j \in \mathbf{A}} f(A^j)$. Thus, we would only need designers to estimate the baseline probability $P(A_0^3 > \mathbf{C})$ for some set of alternatives $\mathbf{C}$ in the prediction domain. Due to space limitations, we leave the details to the reader.

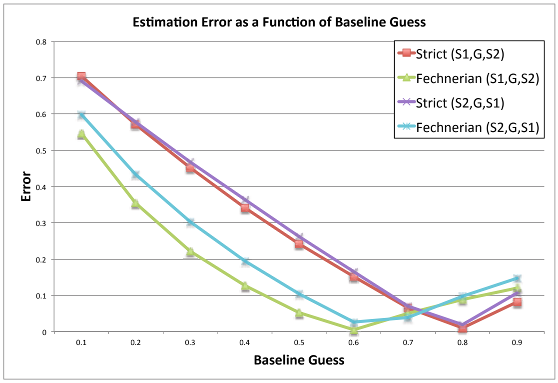

Fig. 1. A plot of the estimation error (absolute difference between observed and predicted $P(A_1 > C_0)$) as a function of baseline guess $P(A_0 > C_0)$. There are series for both the strict and Fechnerian utility models in the S_1, G, S_2 and S_2, G, S_1 settings.

6 Conclusion

In this paper, we have presented a method for transferring probabilistic behavior data from one domain to another using models from mathematical psychology. We have stated the assumptions and conditions under which this approach is reasonable to apply and shown that it can be quite accurate. Using this method, designers will no longer have to hand-author entire models using intuition or expert knowledge, but can rely on lessons learned in other domains. In addition to formally deriving two-choice alternative models, we have described how it can apply more generally to n-choice alternatives.

To characterize the performance of these models, we used data collected from various literatures as well as inputs we varied. The results from using these data and inputs suggest that 1) our models, especially the Fechnerian utility model, can be very accurate under certain conditions; and 2) the largest influence on the overall accuracy of the model is the author-provided baseline estimate for the prediction domain, and not the quality of the source data or the transfer process. What our results do not yet describe—this is a topic for future research—is a concise understanding of the conditions when our models will be most applicable and accurate. Despite this, we are encouraged that this approach will be useful for accurately constructing player models.

References

1. Bates, J.: Virtual Reality, Art, and Entertainment. Presence: The Journal of Teleoperators and Virtual Environments 2, 133–138 (1992)
2. Cialdini, R.B.: Influence: The Psychology of Persuasion. Collins (1998)
3. Falmagne, J.C.: Foundations of Fechnerian Psychophysics. In: Krantz, D.H., Atkinson, R.C., Luce, R.D., Suppes, P. (eds.) Contemporary Developments in Mathematical Psychology. Freeman (1974)
4. Folkes, V.S., Martin, I.M., Gupta, K.: When to Say When: Effects of Supply on Usage. Journal of Consumer Research 20, 467–477 (1993)

5. Guilford, J.P.: Psychometric Methods, 2nd edn. Macmillan (1954)
6. Johnson, E.J., Bellman, S., Lohse, G.: Defaults, framing and privacy: Why opting in $\neq$ opting out. Marketing Letters 13(1), 5–15 (2002)
7. Laurel, B.: Toward the Design of a Computer-Based Interactive Fantasy System. Ph.D. thesis, Drama department, Ohio State University (1986)
8. Luce, R.D.: Individual Choice Behavior. Wiley (1959)
9. Luce, R.D., Galanter, E.: Discrimination. In: Luce, R.D., Bush, R.R., Galanter, E. (eds.) Handbook of Mathematical Psychology, vol. 1, pp. 191–243 (1963)
10. Mateas, M.: An Oz-Centric Review of Interactive Drama and Believable Agents. In: Veloso, M.M., Wooldridge, M.J. (eds.) Artificial Intelligence Today. LNCS (LNAI), vol. 1600, pp. 297–328. Springer, Heidelberg (1999)
11. Nelson, M.J., Mateas, M., Roberts, D.L., Isbell, C.L.: Declarative Optimization-Based Drama Management in the Interactive Fiction Anchorhead. IEEE Computer Graphics and Applications (Special Issue on Interactive Narrative) 26(3), 30–39 (2006)
12. Nelson, M.J., Roberts, D.L., Isbell, C.L., Mateas, M.: Reinforcement Learning for Declarative Optimization-Based Drama Management. In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2006 (2006)
13. Regan, D.T.: Effects of a Favor and Liking on Compliance. J. of Experimental Social Psychology 7, 627–639 (1971)
14. Roberts, D.L.: Computational Approaches for Reasoning About and Shaping Player Experiences in Interactive Narratives. Ph.D. thesis, Georgia Tech. (2010)
15. Roberts, D.L., Bhat, S., Clair, K.S., Isbell, C.L.: Authorial Idioms for Target Distributions in TTD-MDPs. In: Proc. of the 22nd Conf. on Artificial Intelligence, AAAI 2007 (2007)
16. Roberts, D.L., Furst, M.L., Isbell, C.L., Dorn, B.: Using Influence and Persuasion to Shape Player Experiences. In: 2009 Sandbox: ACM SIGGRAPH Video Game Proceedings, SIGGRAPH Sandbox (2009)
17. Roberts, D.L., Isbell, C.L.: A Survey and Qualitative Analysis of Recent Advances in Drama Management. International Transactions on Systems Science and Applications, Special Issue on Agent Based Systems for Human Learning 3(1), 61–75 (2008)
18. Roberts, D.L., Nelson, M.J., Isbell, C.L., Mateas, M., Littman, M.L.: Targeting Specific Distributions of Trajectories in MDPs. In: Proceedings of the 21st National Conference on Artificial Intelligence, AAAI 2006 (2006)
19. Sullivan, A., Chen, S., Mateas, M.: Integrating Drama Management into an Adventure Game. In: Proc. of the Fourth Conf. on Artificial Intelligence and Interactive Digital Entertainment (AIIDE 2008). AAAI Press (2008)
20. Sullivan, A., Chen, S., Mateas, M.: From Abstraction to Reality: Integrating Drama Management into a Playable Game Experience. In: Proc. of the 2009 AAAI Spring Symposium on Intelligent Narrative Technologies II (2009)
21. Thurstone, L.L.: A Law of Comparative Judgement. Psych. Review 34, 273–286 (1927)
22. Thurstone, L.L.: Psychophysical Analysis. American J. of Psychology 38, 368–389 (1927)
23. Weyhrauch, P.: Guiding Interactive Drama. Ph.D. thesis, School of Computer Science (1997)
24. Zermelo, E.: Die berechnung der turnier-ergebnisse als ein maximumproblem der wahrscheinlichkeitsrechnung. Mathematische Zeitschrift 29 (1929)