Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

NEURAL NETWORKS AND AGENT-BASED DIFFUSION MODELS

Ashkan Negahban

School of Graduate Professional Studies
The Pennsylvania State University
30 East Swedesford Road
Malvern, PA 19355, USA

ABSTRACT

This paper introduces a new consumer decision-making model where each agent uses a neural network to evaluate word-of-mouth and predict her utility prior to adoption a new product based on her experiences in the past. The model considers the fact that consumers may not know their true preferences before experiencing the product. By using a neural network, an agent can: (1) interpret the feedback from a neighbor who has conflicting preferences with her; (2) interpret partially positive and/or negative feedback; and, (3) assign different weights to the feedback received from different neighbors. The model is implemented in an agent-based simulation model to verify that the resulting diffusion dynamics follow a typical diffusion curve. Preliminary experiments with the model also provide interesting results about the effect of the number of product attributes on the quality of an individual's utility prediction as well as proportion of satisfied adopters.

1 INTRODUCTION

The consumer marketplace is a complex system, where market dynamics emerge as a result of the decisions and interactions of consumers, investors, manufacturers, vendors, and distributors. The interrelationship between these components add a layer of complexity to marketing research that may not exist in many other disciplines (Rosser 1999). This also means an abundance of opportunity for the application of agent-based modeling and simulation (ABMS). The advantages of ABMS over traditional *aggregate* (population-based) models are as follows: (1) ABMS provides a bottom-up approach that enables modeling population heterogeneity in terms of attributes, preferences, goals, and decision-making behavior; (2) it provides the means to explicitly model consumer social network structure and word-of-mouth operation, and analyze their impact on the market dynamics; and, (3) it has the potential to not only capture emergent market behavior commonly observed in the real world, but also explain the complex non-linear macro-level dynamics by tracing them back to the micro-level behavior and interactions of the model constituents.

Innovation diffusion (a.k.a., new product diffusion) is one of the areas that has benefited substantially from this *relatively* new analysis tool. Early studies on new product diffusion go back to the 1960's (Bass 1969), and numerous aggregate analytical models have been proposed ever since to enhance demand forecasting and decision-making (Bass 2004). However, there are several areas that are beyond the reach of these population-level models that ABMS made possible to explore:

- Targeting and timing of marketing, seeding, and promotional efforts, see Delre et al. (2007), Negahban and Smith (2017), Negahban (2013).
- Impact of population heterogeneity in terms of consumer innovativeness, susceptibility to external (mass media advertisement) and internal (word-of-mouth) influences, adoption behavior, and personal preferences about product attributes (Delre, Jager, and Janssen 2007, Rahmandad and Sterman 2008).

- The role of opinion leaders, i.e., individuals with disproportional influence on others, see Haenlein and Libai (2013), Delre et al. (2010), and van Eck, Jager, and Leeflang (2011).
- Optimal production and sales strategies under various theoretical and empirical social network structures, see Negahban, Yilmaz, and Nall (2014) and Amini et al. (2012).
- Positive and negative word-of-mouth, see Goldenberg et al. (2007) and Goldenberg, Libai, and Muller (2001).

Despite it advantages over traditional aggregate models and its contributions to the innovation diffusion research, the acceptance of ABMS in top academic journals and marketing practice has been slow (Rand and Rust 2011, Leombruni and Richiardi 2005). The findings of several review papers suggest that the existing skepticism toward the application of agent-based simulation in this field can be attributed to the following issues (see, Negahban and Yilmaz (2014), Kiesling et al. (2011), Heath, Hill, and Ciarallo (2009), Zenobia, Weber, and Daim (2009)):

- Lack of clear guidelines on model verification, validation, and calibration;
- Incomplete description of the model and its components in existing ABMS articles;
- Lack of empirical data on and guidelines to estimate the consumer social network structure;
- Simplistic models of consumer decision-making behavior.

The proposed consumer decision-making model in this paper aims at addressing the fourth issue. Consumers generally use rational, intuition, past experience, personal preferences, interpretation of feedback received from their social ties, and their overall impression about a product to make purchasing decisions. However, the existing models generally fail to capture such sophistication involved in human thinking processes. This paper presents the preliminary results of the early development stages of a new consumer decision-making model based on artificial neural networks (ANN) in the hope to address some of the deficiencies in the existing models as identified in the next section. It is worth noting that the idea of incorporating artificial intelligence into agents' behavior, per se, is not new. In fact, several review papers on the application of ABMS in different fields either report or recommend the use of learning mechanisms such as genetic algorithms or neural networks in agent architectures (Gilbert and Terna 2000, Tesfatsion 2002). However, to the best of the author's knowledge, the proposed model in this paper is the first to use artificial neural networks for modeling consumer adoption behavior in the context of an agent-based diffusion model.

The remainder of this paper is organized as follows. Section 2 provides a critical analysis of the commonly used consumer decision-making models in the innovation diffusion literature. Section 3 presents the proposed model and explains how it addresses the identified gaps in the existing models. Model implementation, verification and validation processes, and the results of a set of preliminary experiments with the model are presented in Section 4. Finally, Section 5 outlines opportunities for future research.

2 BACKGROUND: EXISTING CONSUMER DECISION-MAKING MODELS

The majority of the agent-based diffusion studies, including this paper, model the diffusion of a new product or information among a population of individuals (e.g., potential consumer agents) as a contagious epidemic transmitted in a network of humans, and use a variant of the SIS (Susceptible-Infected-Susceptible) or SIR (Susceptible-Infected-Removed) models (Dodds and Watts 2005). In the SIS model, potential consumers are initially *undecided* (susceptible) and may decide to adopt/reject (or simply become aware of) the new product (infected) if they are directly linked with one or more adopter/rejector/informed consumer (i.e., word-of-mouth). When an agent makes her adoption decision, she is not *removed* form the market and may become susceptible again (consider repeat purchases). The SIR model is the same except that once infected, the agent will be removed from the list of *potential* consumers (her decision is final).

What makes these models different is the word-of-mouth mechanism and the rules that govern agents' decision-making process. The existing consumer decision-making models fall under two general categories:

- *Threshold models:* In threshold models, each agent has a personal threshold with respect to one or more types of stimuli. If the magnitude of the stimulus exceeds this threshold, a certain reaction or activity is evoked. In Delre, Jager, and Janssen (2007) and Delre et al. (2007), consumers adopt the new product if it meets a minimum utility level. The utility is a weighted function of personal preference and social influence, that are both threshold functions with binary output. The first component compares the product's quality versus the individual's minimum preference, and the second component compares the percentage of adopters in the individual's personal network with a threshold that determines the consumer's susceptibility to social influence. The model in van Eck, Jager, and Leeflang (2011) is similar except that the social pressure is a continuum (i.e., there is no threshold for social influence). As more neighbors adopt the product, the normative influence in favor of the product increases. In another related model, Delre et al. (2010) consider an adoption threshold for overall utility, while individual preference and social influence are both continuous functions.
- *Probability-based models:* In these models, transitions of potential consumers from one state to another are probabilistic. For instance, the model in Rand and Rust (2011) only considers positive word-of-mouth, and the probability of adoption is a function of the individual's coefficients of innovation and imitation (the same parameters in the model by Bass (1969)) and the proportion of adopter neighbors. In another model that accounts for both positive and negative word-of-mouth, the probability of adoption, rejection, or remaining undecided depend not only on the individual's innovation and imitation coefficients, but also the number of satisfied and dissatisfied adopters in her local network, see Goldenberg et al. (2007), Negahban and Smith (2017), and Amini et al. (2012). the term *probability-based* should not be interpreted as if agent-based diffusion models that use threshold models are not stochastic. The focus here is only on how the status of an agent is updated (e.g., from undecided to adopter). In threshold models, an undecided consumer adopts the new product (with a probability of 1.00) if the utility exceeds her adoption threshold, while in probability-based models the consumer adopts based on the adoption probability.

Some of the limitations of these models can be summarized as follows:

1. Learning from previous interactions with social ties is ignored.
2. In the majority of these models, all neighbors have the same *weight*. Even those studies that distinguish between *strong* and *weak* ties use the same weight for all neighbors in each group (for example, see Goldenberg et al. (2007)).
3. Feedback from a neighbor is either positive or negative (i.e., no partial *membership*). In reality, it is sometimes difficult to clearly distinguish between positive and negative word-of-mouth as consumers' feedback can be ambiguous and/or subjective. Consider the following product review about a certain model of robot vacuums (www.cnet.com, March 7, 2014): "... This brushless technology removes debris better than previous iterations and keeps maintenance to a minimum...The bin is supposedly larger than in previous models, but it still fills up too quickly when you're dealing with dust and shedding pets..." . Comments like this do not neatly fall under positive or negative category.
4. Consider two friends that have completely different preferences about the screen size for a smartphone. When one of them complains that the screen is too big, the net effect on the other individual will most likely be positive. The existing models fail to capture situations like this.
5. While population heterogeneity in terms of innovativeness and imitation level are addressed, the general decision-making logic is assumed to be the same for all consumers. In other words, given a set of inputs, all individuals would make the exact same decision (or at least the probability of

making a certain decision will be the same). In reality, individuals' thought processes are different and people may react differently to the same situation.

6. Most threshold models assume that individuals are fully aware of their preferences before adopting the product. For a new product, it is safe to say that consumers may not know exactly what their preferences are with respect to the different product attributes until they actually experience it. It has been known for a long time that consumers' preferences evolve as they go through different stages of learning before and after trial of the new product (Carpenter and Nakamoto 1989).

The proposed model in this paper falls under the general category of threshold models and strives to address the above issues by incorporating artificial neural networks into the consumer decision-making model. The following section explains in detail the different components of the model.

3 THE PROPOSED CONSUMER DECISION-MAKING MODEL BASED ON ANNs

The adoption status and operation of word-of-mouth is summarized in Figure 1. Consumers interact with each other through word-of-mouth (WOM) and make decisions regarding the adoption of a new product. The model accounts for both innovators and imitators. To model innovators, at any time t, an undecided consumer i adopts with a probability p_i (a.k.a., coefficient of innovation). This is consistent with the assumptions of fundamental diffusion models that innovators' decisions are independent of the social influence (Bass 1969). Imitators, on the other hand, will start evaluating the product once they receive feedback from their direct social ties. Similar to the threshold models described in Section 2, at any time t, an undecided agent i makes decisions based on the utility that she expects from the product ($U_{i,t}^{expected}$). Each agent has an adoption threshold (U_i^{adopt}) and a rejection threshold (U_i^{reject}) such that $U_i^{reject} \leq U_i^{adopt}$. The agent's decision is then determined as follows:

$$Decision = \begin{cases} adopt & \text{if } U_{i,t}^{expected} \geq U_i^{adopt}, \\ wait & \text{if } U_i^{reject} < U_{i,t}^{expected} < U_i^{adopt}, \\ reject & \text{if } U_{i,t}^{expected} \leq U_i^{reject}. \end{cases} \tag{1}$$

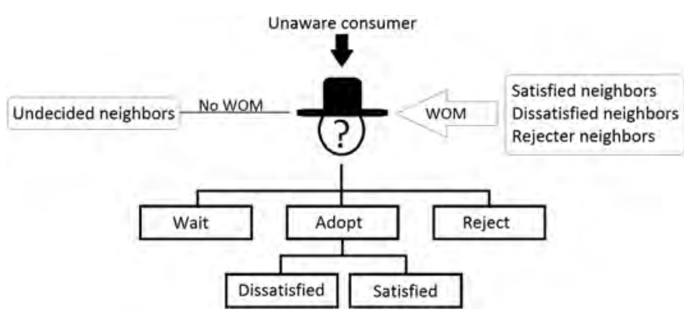

Figure 1: Adoption status and the operation of word-of-mouth.

Since rejectors never experience the product, they communicate their *expected* utility to their social ties. An adopter, on the other hand, will share her actual experience with the product (U_i^{actual}). Undecided/waiting customers do not participate in the word-of-mouth operation. Consider a product with n attributes $a_1, a_2, ..., a_n$. In version of the model, the actual utility is a function of the *normalized* euclidean distance between the product's specifications and the individual's preferences given by

$$U_i^{actual} = 1 - \sqrt{\frac{(\pi_1 - a_1)^2 + (\pi_2 - a_2)^2 + ... + (\pi_n - a_n)^2}{n}}, \tag{2}$$

where π_j is the individual's preference for attribute j. Adopters will either become *satisfied* or *dissatisfied* with the product by comparing their actual utility with their satisfaction threshold as follows:

$$Decision = \begin{cases} satisfied, & U_i^{actual} \geq U_i^{satisfaction}, \\ disatisfied, & otherwise. \end{cases} \quad (3)$$

A potential consumer will become aware of the product only if there is at least one rejector or adopter in her local network. The magnitude of social pressure increases when the potential consumer receives feedback from more neighbors (Granovetter and Soong 1986). As the social pressure increases, it is more likely that the undecided consumer will consider evaluating the product. The probability ρ that an undecided agent will evaluate the product is directly related to the proportion of WOM contributers in her local network:

$$\rho = \frac{Number\ of\ WOM\ contributing\ neighbors}{Number\ of\ neighbors}. \quad (4)$$

The individual's expected utility at time t ($U_{i,t}^{expected}$) is a weighted function of two components, namely *initial perception* and *social influence*, and is given by

$$U_{i,t}^{expected} = (1-\beta_i)y_i + \beta_i O_{i,t}, \quad (5)$$

where y_i is the effect of the individual's *perception* about the product prior to adoption, $O_{i,t}$ is the social influence at time t, and β_i determines the relative importance of the social influence ($0 \leq \beta_i \leq 1$). Similar to Delre, Jager, and Janssen (2007), Delre et al. (2007), and van Eck, Jager, and Leeflang (2011), y_i is considered to be a threshold function in this version of the model, we have

$$y_i = \begin{cases} 1 & \gamma_i \geq U_i^{adopt}, \\ 0 & otherwise, \end{cases} \quad (6)$$

where γ_i is the individual's initial perception of the product before adoption. The social influence component ($O_{i,t}$), on the other hand, is based on the feedback from direct neighbors. The modeling approach for this component is the main part that differentiates this model from existing models.

3.1 Artificial Neural Networks for Evaluating Social Influence

The undecided consumer evaluates the social influence based on past experiences, i.e., based on her neighbors' feedback about similar products in the past and her actual experience with those products. The consumer agent learns how to interpret the feedback from her social ties and predicts her utility for the new product through a neural network (Figure 2).

As discussed previously, an undecided consumer is aware of the actual utility of adopter neighbors and expected utility of those neighbors that rejected the product. Consider an agent with l direct ties. The input into the agent's ANN at time t is a vector of size l as follows:

$$[I_{1t}, I_{2t}, \ldots, I_{lt}],$$

where, for each neighbor $k = 1, 2, \ldots, l$, we have

$$I_{kt} = \begin{cases} U_k^{actual} & \text{if } neighbor\ k\ is\ an\ adopter, \\ U_{k,\tau}^{expected} & \text{if } neighbor\ k\ rejected\ the\ product\ at\ time\ \tau < t, \\ 0 & \text{if } neighbor\ k\ is\ undecided\ or\ unaware\ (not\ a\ WOM\ contributor). \end{cases}$$

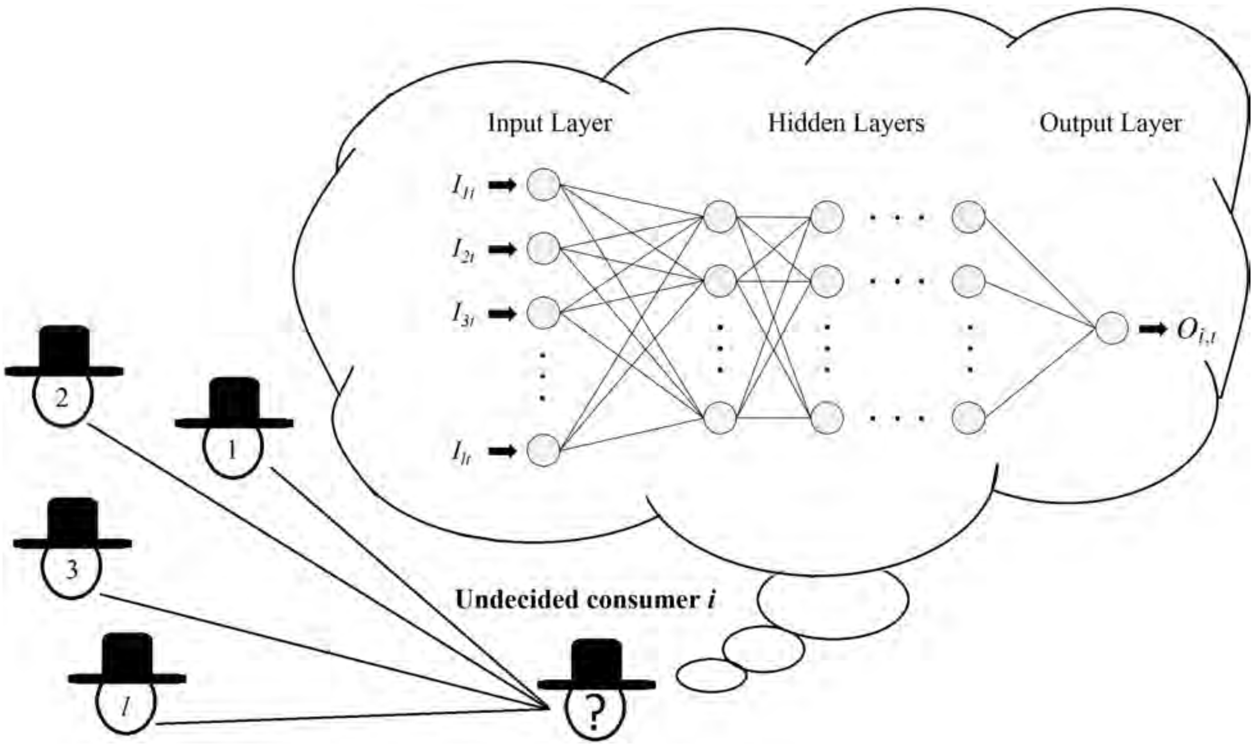

Figure 2: The general structure of a consumer agent's multi-layer feed-forward neural network. The net may or may not include *biases* (i.e., neurons that always take a value of 1).

The neural network predicts the utility of the agent for the new product, and thus the output is a scalar between 0 and 1 since $O_{i,t} = \widehat{U}_i^{actual}$. In other words, consumer i's utility at any time t is assumed to be a function (perhaps a highly nonlinear and complex function) of her neighbors' decision:

$$U_i^{actual} = f_t(I_{1t}, I_{2t}, ..., I_{lt}). \tag{7}$$

The structure of this function is unknown in general since the consumer is only aware of her neighbors' adoption decision and their *overall* utility, while she does not know their exact preferences for the different product attributes. This makes a neural network an appropriate tool to approximate this function. Before discussing the implementation of the proposed consumer decision-making model and other relevant aspects of the simulation process (such as ANN structure and training, social network structure and size, etc.), it is important to understand how this model addresses the drawbacks of existing models.

3.2 Advantages of the ANN-Based Consumer Decision-Making Model

The proposed model has several advantages over existing models (the following items have a one-to-one association to the gaps identified in Section 2):

1. Incorporating a neural network into the consumer agents' architecture allows for learning from past experience. More specifically, consumers evaluate the word-of-mouth received and predict their utility for the new product based on their experience with similar products and feedback from their social ties about those products in the past.
2. In theory, a multi-layer neural network (with one or more hidden layers) can learn any continuous function to an arbitrary accuracy. The *importance* given to each neighbor is determined by the outcome of the training process, which may not necessarily lead to equal weights for all neighbors.

3. The net's inputs (i.e., neighbors' utilities) can take any value between 0 and 1, and thus the proposed model allows for any kind of feedback from 100% positive (1) to 100% negative (0).
4. Consider a consumer agent and one of it's neighbors with completely different preferences. If the neighbor's utility for a product is close to 1, the consumer's utility is expected to be close to zero. The consumer's neural net can learn such patterns in the training data set so that a negative feedback (low utility) from this neighbor would increase the social influence in favor of the product.
5. The proposed model model not only allows for population heterogeneity in terms of preferences, adoption/rejection/satisfaction thresholds, susceptibility to social influence (β_i), and initial perception about the product, but also allows for heterogeneity in terms of the decision-making logic. Consumers' neural nets can have different number of hidden layers and/or hidden neurons in each layer. Even if the same general neural net structure is use for all consumers, the networks will be trained differently since the inputs are associated with each consumer's neighbors. Unless the social network is a lattice, the number of neighbors and thus input neurons will also be different.
6. The proposed model accounts for the fact that a consumer may not be fully aware of her preferences regarding different product attributes before using the product. Note that the parameter γ_i and consequently y_i are based on an *overall* perception about the product prior to adoption.

4 MODEL IMPLEMENTATION

An agent-based model is developed that implements the proposed ANN-based consumer model. The model is developed in Repast Simphony (North et al. 2013). Parameters choices are summarized in Table 1.

The market consists of 1600 consumer agents connected through a small-world network (Watts and Strogatz 1998). A population of size of 1000 is found to be sufficiently large to effectively capture the diffusion dynamics and provide statistically valid results (Cowan and Jonard 2004). The small-world network is generated by randomly rewiring some of the links in a lattice (a lattice with a degree of 12 is used to generate the network). Initially, all consumers are unaware of the new product. The product is released at $t = 0$, and at every time tick, consumers make their adoption decision based on the proposed model. The simulation run ends when the market is exhausted.

4.1 Consumers' Neural Nets and Training Process

Consumers' neural nets are implemented using the Java Object Oriented Neural Engine (Joone) library (http://www.jooneworld.com/). The structure of the neural nets and other parameters related to the training process are chosen based on the guidelines in Fausett (1994). As discussed in Section 3.2, the number of input neurons is equal to the size of the consumer's local network. A single hidden (sigmoid) layer is used for all consumers. Theoretically, this is sufficient to capture any mapping to an arbitrary accuracy. As a rule-of-thumb, $\sqrt{l}$ would be an appropriate number of hidden neurons for a neural net with l input neurons. In this model, the number of hidden neurons is set to 5 since some agents may end up with more than 12 neighbors after rewiring the links when generating the small-world network.

Before releasing the product, consumers' neural networks are trained. Specifications for the set of training products are randomly generated. As a rule-of-thumb, for a net to be trained to correctly predict approximately $100(1-e/2)$ of the training set, and $100(1-e)$ of the test set, and given w weights to be estimated, w/e would be an appropriate number of training products. In this model, an average of $12 \times 5 + 5 = 65$ weights need to be estimated for each neural net. Given an acceptable error of $e = 0.15$, $65/0.15 = 433.33$ can be considered as an appropriate size for the training set. Here, 450 training products are used. Training stops after 5000 epochs or if the root-mean-square error (RMSE) drops below 0.0005, whichever occurs first. It is worth noting that the goal here is not to find the *optimal* neural network for each agent, but rather develop a net that can provide reasonably accurate predictions.

Table 1: Parameter choices for the agent-based simulation model.

Parameter	Value/Range	Selected source
Consumer agents		
Population size	1600	Cowan and Jonard (2004)
Coefficient of innovation (p_i)	$U(0.01, 0.05)$	Negahban, Yilmaz, and Nall (2014)
Social influence coefficient (β_i)	$N(\mu_\beta, \sigma_\beta)$	Delre et al. (2007)
μ_β	0.1-0.9 in 0.1 increments	Delre et al. (2007)
σ_β	0.01	Delre et al. (2007)
Initial perception (γ_i)	$U(0,1)$	van Eck, Jager, and Leeflang (2011)
Adoption threshold (U_i^{adopt})	$U(0,1)$	Delre et al. (2007)
Rejection threshold (U_i^{reject})	$U(0, U_i^{adopt})$	None (new feature)[1]
Satisfaction threshold ($U_i^{satisfaction}$)	$U(0,1)$	None (new feature)[1]
Preferences and product attributes		
Number of product attributes (n)	1, 2, 3, and 4	Delre et al. (2007)[2]
Consumer preference for attribute j (π_j)	$U(0,1)$	Delre et al. (2007)[2]
Product specification for attribute j (a_j)	$U(0,1)$	Delre et al. (2007)[2]
Consumer's neural network[3]		
Net type	Multi-layer feed-forward	Fausett (1994)
Number of input neurons	Size of local network	N/A
Number of hidden layers	1	Fausett (1994)
Number of hidden neurons	5	Fausett (1994)
Size of the training set	450	Fausett (1994)
Number of epochs (training iterations)	5000	Fausett (1994)
The social network		
Network type	Small-world	Watts and Strogatz (1998)
Small-world degree	12	Goldenberg et al. (2007)
Small-world rewiring probability	0.01	Watts and Strogatz (1998)

[1] Existing threshold models do not consider rejection or satisfaction.

[2] Existing threshold models consider only one attribute.

[3] Since this is the first model that uses an ANN-based adoption model, parameter choices under this category are based on general guidelines on designing and training neural networks.

4.2 Verification and Validation

Before performing the experiments, verification and validation (V&V) were performed to gain confidence about the correct implementation and feasibility of the model. For the sake of conciseness, a brief description of the V&V steps is provided here.

The four aspects of validation for agent-based models in the marketing domain as suggested by Rand and Rust (2011) are assessed as follows. At the *micro-face level*, consumer agents possess a local social network that prevents them from obtaining information about the entire population. Individuals' decisions are driven by their innovativeness (p_i), initial perception (γ_i), and susceptibility to social influence (β_i), which capture the effect of mass media advertisement and word-of-mouth. Population heterogeneity is also considered as consumers have different threshold values, initial perceptions, and neural nets for interpreting word-of-mouth. The underlying mechanism to determine consumers' utility is in line with fundamental theories on consumer behavior (Granovetter and Soong 1986). At the *macro-face level*, the underlying theory of the model is inline with well-established and commonly used diffusion models in practice (such as Bass (1969)). Parameter choices are based on theoretical and/or empirical studies on diffusion dynamics and neural networks, strengthening validity at the *empirical input level*. As for *empirical output validation*, aggregate results follow typical S-shaped diffusion patterns observed in the real world.

Model verification is performed by following the guidelines in Sargent (2005), Macal and North (2010), and Balci (1994). For instance, structured walk-throughs and tracing are performed to verify correct implementation of agents activation and scheduling of methods at each time step, random number generation and probabilistic decisions, synchronous updating of consumers' adoption status, and accessing agents' personal social network. Operational graphics are used to observe the diffusion dynamics over time and verify the model's behavior at the population level. Degenerate tests are also performed to verify the behavior of the model in extreme conditions such as coefficient of innovation of 0 and 1.

4.3 Preliminary Findings

The preliminary experiments explore the effect of two important parameters, namely the relative importance of the social influence (β_i) and the number of product attributes (n). It is important to note that the goal here is not to provide a comprehensive analysis on different model parameters, but rather to test the feasibility and applicability of the proposed ANN-based consumer decision-making model. Based on the analysis of the confidence interval half-widths from a set of pilot runs, the number of replications is set to 20 for each other 36 scenarios. Therefore, more than 1.1 million neural net are trained in total which require approximately 4,000 hours (166 days) of computational time (less than a month in real time as the experiments were run in parallel on a machine with 3.40GHz Core i7 CPU and 32GB of RAM).

Figure 3(a) shows the average demand for different values of the mean social influence coefficient (μ_β). The demand patterns generated by the model seem to follow the typical diffusion dynamics observed in many real-world situations, which affirms that the model has the potential to provide a good fit for real data if calibrated properly. The effect of μ_β can be explained as follows. When μ_β is small, consumers would give much more importance to their initial perception about the product than word-of-mouth. Due to the use of a threshold function to model individuals' perception prior to adoption, even early in the diffusion process, it is highly likely that a consumer will adopt the product simply if her initial perception exceeds her adoption threshold, i.e., if $\gamma_i \geq U_i^{adopt}$. Note that in this implementation of the model, for any individual i, both γ_i and U_i^{adopt} are sampled from a uniform distribution between 0 and 1, and there is a 50% chance that y_i takes a value of 1 (see Equation 5). The net effect will be a fast adoption rate. As μ_β increases, more importance will be given to the social influence and consumers are more likely to wait until they have a highly positive feeling about the product based on the feedback from their social ties before they make an adoption decision. This results in a slower take-off and not only delays the peak demand, but also leads to a smaller peak demand (this does not necessarily mean smaller cumulative adoptions at the end of the diffusion process).

Perhaps the more interesting finding pertains to the effect of the number of product attributes (n). As shown in Figure 3(b), we observe a general decrease in the proportion of satisfied adopters as n increases. This is mainly due to the increased complexity of the function in (7) that consumers' neural nets try to approximate. With only one attribute, the neural net can easily identify a neighbor with very similar or very different preferences compared to the parent consumer, allowing the net to make much more accurate predictions. As the number of product attributes increases, the mapping of neighbors' utilities becomes much more complex, increasing the chance of misinterpreting the social influence and dissatisfaction. It is important to note that the data include innovators (individuals that adopt the product randomly, i.e., independent of the social pressure). The satisfaction status/outcome for these adopters is random. The effect of the number of product attributes on satisfaction probability would be even more obvious if we only consider imitator adopters (i.e., individuals that evaluate the word-of-mouth using their neural network).

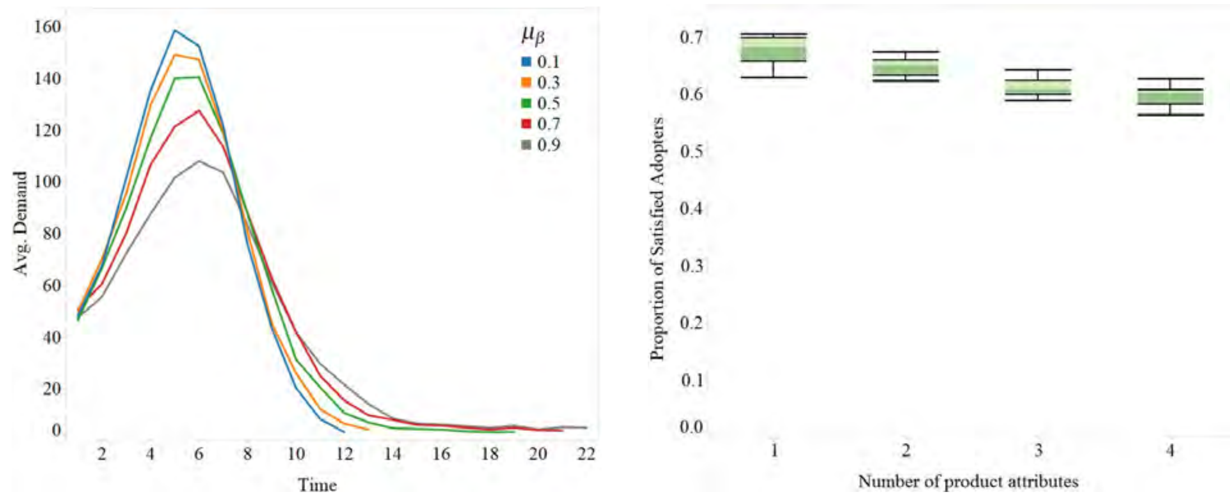

(a) The effect of mean social influence strength (μ_β). The number of product attributes is 4.

(b) The effect of the number of product attributes (n). The mean social influence coefficient is 0.7.

Figure 3: The results of preliminary experiments with the ANN-based adoption behavior model.

5 CONCLUSIONS

A new consumer decision-making model based on neural networks is proposed that allows for aggregating and interpreting word-of-mouth based on past experience. The advantages of the model over existing approaches are discussed. The model is then implemented in an agent-based simulation model of innovation diffusion. Preliminary analyses indicate the applicability of the model and confirm that it can generate diffusion dynamics similar to those observed in the real world. The results also provide interesting insights on the relationship between the number of product attributes and proportion of satisfied consumers.

The proposed model unfolds an abundance of future research opportunities. Perhaps the biggest challenge is the computational effort required to train thousands of neural networks in every replication of the simulation. There is a need for efficient methods that allow distributing not simulation runs but *agents* across multiple processors so that multiple nets can be trained in parallel. Potential extensions/variations of the model may consider: (1) availability of partial information on neighbors' preferences about specific product attributes; (2) exploring population heterogeneity in terms of learning capabilities by using various types of neural networks (e.g., single-layer vs multi-layer); (3) alternative methods for calculating the similarity between consumer preferences and product attributes; and, (4) continuous learning as opposed to the current implementation of the model where neural nets are fixed after training is complete.

Parameter estimation and calibration is another important area for future research. Several methodologies have been proposed to identify product features and associated consumer opinions from large-scale social media data (Tuarob and Tucker 2015). These methods can be extended to collect data in order to estimate parameters related to consumer interactions, social network structure, and the effect of word-of-mouth. The effect of overfitting, a common problem in training neural nets, also needs to be addressed using any of the existing methods such as the one presented in Srivastava et al. (2014). Finally, while originally developed for the innovation diffusion context, general agent architectures can be developed for other application areas such as social sciences, manufacturing, healthcare (Powell and Mustafee 2016), and military (Song, Zhang, and Qian 2013).

REFERENCES

Amini, M., T. Wakolbinger, M. Racer, and M. G. Nejad. 2012. "Alternative Supply Chain Production-Sales Policies for New Product Diffusion: An Agent-Based Modeling and Simulation Approach". *European Journal of Operational Research* 216 (2): 301–311.

Balci, O. 1994. "Validation, Verification, and Testing Techniques throughout the Life Cycle of a Simulation Study". *Annals of Operations Research* 53 (1): 121–173.

Bass, F. M. 1969. "A New Product Growth for Model Consumer Durables". *Management Science* 15 (5): 215–227.

Bass, F. M. 2004. "Comments on "A New Product Growth for Model Consumer Durables"". *Management Science* 50 (12): 1833–1840.

Carpenter, G. S., and K. Nakamoto. 1989. "Consumer Preference Formation and Pioneering Advantage". *Journal of Marketing Research* 26 (3): 285–298.

Cowan, R., and N. Jonard. 2004. "Network Structure and the Diffusion of Knowledge". *Journal of Economic Dynamics and Control* 28 (8): 1557–1575.

Delre, S., W. Jager, T. Bijmolt, and M. Janssen. 2007. "Targeting and Timing Promotional Activities: An Agent-Based Model for the Takeoff of New Products". *Journal of Business Research* 60 (8): 826–835.

Delre, S., W. Jager, and M. Janssen. 2007. "Diffusion Dynamics in Small-World Networks with Heterogeneous Consumers". *Computational & Mathematical Organization Theory* 13:185 – 202.

Delre, S. A., W. Jager, T. H. A. Bijmolt, and M. A. Janssen. 2010. "Will It Spread or Not? The Effects of Social Influences and Network Topology on Innovation Diffusion". *Journal of Product Innovation Management* 27 (2): 267–282.

Dodds, P., and D. Watts. 2005. "A Generalized Model of Social and Biological Contagion". *Journal of Theoretical Biology* 232 (4): 587 – 604.

Fausett, L. 1994. *Fundamentals of Neural Networks: Architectures, Algorithms and Applications*. 1^{st} ed. Upper Saddle River, NJ, USA: Prentice-Hall, Inc.

Gilbert, N., and P. Terna. 2000. "How to Build and Use Agent-Based Models in Social Science". *Mind & Society* 1 (1): 57 – 72.

Goldenberg, J., B. Libai, S. Moldovan, and E. Muller. 2007. "The NPV of Bad News". *International Journal of Research in Marketing* 24 (3): 186–200.

Goldenberg, J., B. Libai, and E. Muller. 2001. "Talk of the Network: A Complex Systems Look at the Underlying Process of Word-of-Mouth". *Marketing Letters* 12:211 – 223.

Granovetter, M., and R. Soong. 1986. "Threshold Models of Interpersonal Effects in Consumer Demand". *Journal of Economic Behavior & Organization* 7 (1): 83 – 99.

Haenlein, M., and B. Libai. 2013. "Targeting Revenue Leaders for a New Product". *Journal of Marketing* 77 (3): 65–80.

Heath, B., R. Hill, and F. Ciarallo. 2009. "A Survey of Agent-Based Modeling Practices (January 1998 to July 2008)". *Journal of Artificial Societies and Social Simulation* 12 (4): 9 – 43.

Kiesling, E., M. Gunther, C. Stummer, and L. Wakolbinger. 2011. "Agent-Based Simulation of Innovation Diffusion: A Review". *Central European Journal of Operations Research* 20 (2): 183–230.

Leombruni, R., and M. Richiardi. 2005. "Why Are Economists Sceptical about Agent-Based Simulations?". *Physica A: Statistical Mechanics and its Applications* 355 (1): 103 – 109.

Macal, C., and M. North. 2010. "Tutorial on Agent-Based Modelling and Simulation". *Journal of Simulation* 4 (3): 151–162.

Negahban, A. 2013. "A Hybrid Simulation Framework for the Newsvendor Problem with Advertising and Viral Marketing". In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 1613–1624. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Negahban, A., and J. S. Smith. 2017. "A Joint Analysis of Production and Seeding Strategies for New Products: An Agent-Based Simulation Approach". *Annals of Operations Research*.

Negahban, A., and L. Yilmaz. 2014. “Agent-Based Simulation Applications in Marketing Research: An Integrated Review”. *Journal of Simulation* 8 (2): 129–142.

Negahban, A., L. Yilmaz, and T. Nall. 2014. “Managing Production Level in New Product Diffusion: An Agent-Based Simulation Approach”. *International Journal of Production Research* 52 (17): 4950–4966.

North, M. J., N. T. Collier, J. Ozik, E. R. Tatara, C. M. Macal, M. Bragen, and P. Sydelko. 2013. “Complex Adaptive Systems Modeling with Repast Simphony”. *Complex Adaptive Systems Modeling* 1 (1): 3.

Powell, J. H., and N. Mustafee. 2016. “Widening Requirements Capture with Soft Methods: An Investigation of Hybrid M&S Studies in Health Care”. *Journal of the Operational Research Society*.

Rahmandad, H., and J. Sterman. 2008. “Heterogeneity and Network Structure in the Dynamics of Diffusion: Comparing Agent-Based and Differential Equation Models”. *Management Science* 54 (5): 998 – 1014.

Rand, W., and R. T. Rust. 2011. “Agent-Based Modeling in Marketing: Guidelines for Rigor”. *International Journal of Research in Marketing* 28 (3): 181–193.

Rosser, Jr., J. B. 1999. “On the Complexities of Complex Economic Dynamics”. *The Journal of Economic Perspectives* 13 (4): 169 – 192.

Sargent, R. G. 2005. “Verification and Validation of Simulation Models”. In *Proceedings of the 37th Winter simulation Conference*, edited by B. J. A. D. M. D. Rossetti, R. R. Hill and R. G. Ingalls, 130–143. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Song, X., S. Zhang, and L. Qian. 2013. “Opinion Dynamics in Networked Command and Control Organizations”. *Physica A: Statistical Mechanics and its Applications* 392 (20): 5206 – 5217.

Srivastava, N., G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. 2014. “Dropout: A Simple Way to Prevent Neural Networks from Overfitting”. *Journal of Machine Learning Research* 15 (1): 1929 – 1958.

Tesfatsion, L. 2002. “Agent-Based Computational Economics: Growing Economies From the Bottom Up”. *Artificial Life* 8 (1): 55 – 82.

Tuarob, S., and C. Tucker. 2015. “Quantifying Product Favorability and Extracting Notable Product Features Using Large Scale Social Media Data”. *Journal of Computing and Information Science in Engineering* 15 (3): 031003–031003–12.

van Eck, P. S., W. Jager, and P. S. H. Leeflang. 2011. “Opinion Leaders’ Role in Innovation Diffusion: A Simulation Study”. *Journal of Product Innovation Management* 28 (2): 187–203.

Watts, D. J., and S. H. Strogatz. 1998. “Collective Dynamics of ‘Small-World’ Networks”. *Nature* 393 (6684): 440–442.

Zenobia, B., C. Weber, and T. Daim. 2009. “Artificial Markets: A Review and Assessment of a New Venue for Innovation Research”. *Technovation* 29 (5): 338–350.

AUTHOR BIOGRAPHIES

ASHKAN NEGAHBAN is an Assistant Professor of Engineering Management at The Pennsylvania State University, School of Graduate Professional Studies. He has a PhD and MS from Auburn University and a BS from University of Tehran (all in Industrial Engineering). His research interest is in the application of agent-based, discrete event, and Monte Carlo simulation in manufacturing systems and operations management. His email and web addresses are anegahban@psu.edu and http://ashkannegahban.com.