

27: Using Agent-Based Models to Understand Health-Related Social Norms

Gita Sukthankar and Rahmatollah Beheshti

Abstract

Social norms have been demonstrated to strongly affect people's health choices, yet they are often not included in health models due to the complex interdependencies of reasoning about norm adoption over a large population. This article introduces two agent-based modeling frameworks designed explicitly for reasoning about the influence of social norms, Lightweight Normative Architecture (LNA) and Cognitive Social Learners (CSL), and illustrates their usage for modeling smoking cessation trends. LNA models the impact of personal, social, and environmental factors on recognition, adoption, and compliance with a single smoking norm, whereas CSL is capable of reasoning about multiple social norms. By incorporating a more complex normative reasoning model, CSL can not only predict smoking trends but also accurately forecasts population-level responses to surveys on the social acceptability of smoking. These social models are an important complement to existing biological models of human health and wellness.

Introduction

Agent-based models have been shown to be valuable for many types of social simulation problems, including predicting the effects of geography, economic fluctuations, and public policy decisions on human populations. Understanding the influence of social norms on human behavior is an important aspect of performing accurate population-level modeling, and forecasting *norm emergence* has attracted research attention in both the agent-based social simulation and multi-agent system communities.

In this chapter, we describe an agent-based simulation that we constructed to model smoking cessation trends at University of Central Florida following the initiation of a smoke-free campus policy. Since social norms have been shown to strongly affect health-related habits such as overeating, binge drinking, and smoking, our simulation focuses on the social, rather than addictive, elements of the smoking cessation problem. Our Lightweight Normative Architecture (LNA) (Beheshti, Sukthankar, 2014b) models the impact of personal, social, and environmental factors on recognition, adoption, and compliance with campus smoking norms. When initialized with student survey data, it accurately predicts trends in smoking reduction over a one year timeframe.

One weakness with LNA is that it has a relatively simple internal model of the human decision-making process. To address this issue, we created a general normative architecture, Cognitive Social Learners (CSL) (Beheshti et al., 2015), which is capable of reasoning about any social norm. CSL provides a computational mechanism for transitioning behaviors learned during repeated social interactions into the agent's internal cognitive model of preexisting beliefs, desires, and intentions. By incorporating a more complex normative reasoning model, CSL not only predicts smoking trends but also accurately forecasts population-level perception on the social acceptability of smoking.

Related Work

Our models in this chapter examine smoking behaviors from a normative point of view. Non-normative models of smoking behavior already exist; for instance, *SimSmoke* is one of the widely used tobacco control policy simulations. It models the dynamics of smoking use and smoking-attributed deaths in the society of interest, as well as the effects of policies on those outcomes (Levy et al., 2005). Other types of simulations have been used to model the consequences of second-hand smoking (Dacunto et al., 2013). In addition to norms, our proposed approach also simulates network effects as was done in Beckman et al.'s (2011) study on the propagation of adolescent smoking behavior.

Most existing models within the medical and public health community are based on statistical analysis of smoking data (Luo et al., 2015). These methods often focus on a narrow aspect of the problem, such as modeling abstinence due to changes in brain cells. However, some models in the public health domain have been based on system dynamics approaches (Timms et al., 2012). An introduction to this set of techniques can be seen in Homer, Hirsch (2006).

The relationship between social norms and smoking behavior was examined as part of a European Union study on the impact of cultural differences on the emer-

gence of norms in different countries after the commencement of anti-smoking legislation (Dechesne et al., 2013). Our current ABM does not attempt to recreate cultural effects. Rather than studying smoking cessation behavior at the macroscopic level, we adopt a higher fidelity approach in which the daily behavior patterns of individual agents are simulated within an activity-oriented microsimulation.

Lightweight Normative Architecture (LNA)

To construct a normative model for a real-world scenario, we need to define both a norm architecture and the components that are used to recreate the real-world problem.

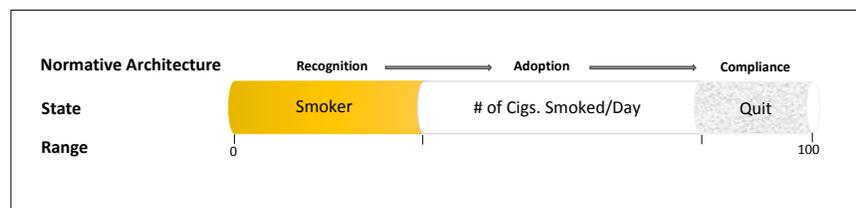


Figure 27.1: A schematic representation of the LNA architecture.

Each agent has a personal *smoking-value* ranging from 0 to 100 that governs its behavior. As shown in Figure 27.1, our architecture contains three stages: **recognition**, **adoption** and **compliance**. In the first stage (recognition), the beliefs of an agent change and develop. During the adoption phase, the agent commences action. Note that the general definition of adoption in normative systems is very consistent with our smoking scenario. During the adoption phase the agent can opt to violate the norm. The equivalent violation in the smoking scenario (recidivism) is quite common in those trying to quit. To quit smoking, a smoker usually decreases the number of smoked cigarettes, which can be considered as another adoption behavior. The compliance phase is used to simulate the time period when the agent seriously attempts to quit smoking.

Cognitive Social Learners Architecture (CSL)

CSL adds an advanced structure of learning and reasoning about norms to the simple decision making process of the LNA architecture. Figure 27.2 shows a schematic view of CSL. In this architecture, the belief, desire and intention components implement the cognitive aspects of norm formation, while the game theoretic (GT) interaction and reinforcement learning (RL) recognition parts implement the social aspects. CSL only

models the rational and social elements of human health-related decisions; it meant to be a complement to existing biological models of craving and dependency (Gutkin et al., 2006).

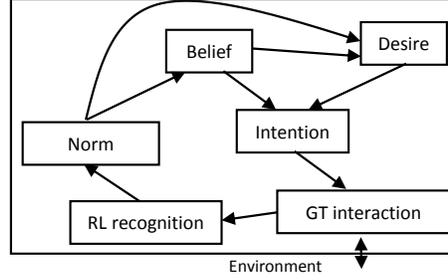


Figure 27.2: Cognitive Social Learners (CSL) Architecture

The representation used for the BDI components and the norms is based on a simplified version of the framework introduced by Casali et al. (2008) and Criado et al. (2010b) in which a certainty degree is assigned to each representation. For example, $(D^- \text{payfine}, 0.45)$ designates a negative desire toward paying a fine with a certainty degree of 0.45.

Belief, Desire, and Intention

The CSL architecture follows a classic BDI structure. Like many normative architectures, each agent is initialized with a set of personal values that model innate preferences. In CSL, these personal values are used to create Type 1 beliefs that have a certainty equal to 1; for instance $(B[\text{hedonism} = 50], 1)$ indicates that the personal value of the agent regarding hedonism is equal to 50. The second form of belief, Type 2, is used to model the agent’s actions, and is represented as $(B[\alpha]\varphi, \delta)$. For instance, $(B[\text{smoking}] \text{bother Rest}, 0.30)$ indicates that the agent believes, with certainty of 0.30, that smoking would bother the other agents.

Desires can be determined independently or based on the agent’s beliefs. Desires are represented as $(D^*\varphi, \delta)$, which models the positive or negative $(* = \{-, +\})$ desire of an agent regarding state φ with certainty of δ . An agent may update its desires when its beliefs change. This process is shown in Equation 27.1; the certainty value of desire D is updated based on function f , which is a user-defined function.

$$((D^*\varphi, \delta_\varphi), (B[\alpha]\varphi, \delta_\phi)) \Rightarrow (D^*\varphi, f(\delta_\varphi, \delta_\phi)) \quad (27.1)$$

Intentions are derived from the set of positive desires, if they have a certainty value

higher than sum of the certainty values of all negative desires relevant to the intention. Equation 27.2 shows this:

$$\begin{aligned} & ((D^+ \varphi_{i_1}, \delta_{\varphi_{i_1}}), \dots, (D^+ \varphi_{i_n}, \delta_{\varphi_{i_n}}), (plan_j, \delta_j)) \\ & \Rightarrow (I_k, f(\delta_{i_1} \dots \delta_{i_n}, \delta_j)) \end{aligned} \tag{27.2}$$

while $\Sigma(\delta_{i_1} \dots \delta_{i_n}) \geq \Sigma(\delta_{l_1} \dots \delta_{l_n})$ and l_1 to l_n are indices of negative desires toward effects of I_k . According to this formula, the set of positive desires (from i_1 to i_n) and plan j will determine the intention k based on a user defined function f . In the smoking case, an agent might have positive desires toward higher happiness and conforming with smoker friends, but negative desires toward becoming ill and being observed by others. In this case, if the sum of certainty values for happiness and consistency is more than the sum of certainty values for becoming sick and being observed (assuming that smoking is part of the agent's current plan), the agent will smoke.

Game-theoretic Interaction

Instead of deciding its actions based on intentions alone, which is often the case in BDI-based methods, the agent's final action is determined after playing a social dilemma game with (one of) the other agents. The other agent could be a neighbor agent or a friend. The maximum certainty value of available intentions is used to create a two-by-two matrix. The two possible actions are performing or refraining from that action. After calculating the payoff value for an action based on the related intentions, fixed values of α and β are used to increase the value of the elements in the matrices representing coordinated action (the agent and its neighbor selecting the same actions) (Easley, Kleinberg, 2010). An example of this matrix for the smoking scenario is shown in Table 27.1. ψ shows the computed payoff value for smoking. ψ' is the payoff for not smoking.

	S	NS
S	$\psi + \alpha$	ψ
NS	ψ'	$\psi' + \beta$

Table 27.1: Example payoff matrix for smoking (S=Smoke, NS=not smoke).

Based on the outcome of games played with this payoff matrix, an agent decides what action to perform. What an agent observes after performing an action may cause an agent to update its personal values (Type 1 beliefs) and learned norms, which in turn modifies its behavior in subsequent steps. For instance, in the case of our example

scenario, after smoking, if there is an advertisement in its vicinity, its value for the effect of advertisements will increase because of the formed habit.

Norm Recognition using RL

The goal of this component is to construct a practical way of recognizing/learning norms, while connecting different components of the architecture. Our RL based recognition component plays the role of a hub among norms and personal values (beliefs) on one hand and the game theoretic interaction on the other hand.

The combination of GT interaction and RL based recognition components is used to implement the social learning process, which propagates norms across the agent population. The aim of the social learning framework is different from similar processes in the domain of multi-agent reinforcement learning. In those the agents play iterative games to learn a policy resulting in a competitive or cooperative equilibrium. Sen and Airiau (2007) note several differences between social learning and multi-agent RL, including the lack of equilibrium guarantees. At every timestep, each agent interacts with a single changing agent, selected at random, from the population. The payoff received by the CSL agent depends only on this interaction. We use a basic Q-learning algorithm for recognizing norms in which states are the discretized current values of an agent's payoff matrices. Learning results in modifications to the certainty degree of available norms. Rewards are calculated based on the changes in the personal values.

Norms

The process of recognizing a social norm is modeled by an agent increasing the norm's certainty value to a positive value. The agent updates the certainty values of norms based on its observations after performing an action. Our norms are represented using the format introduced in Criado et al. (2013), $\langle \Delta, C, A, E, S, R \rangle$, in which Δ designates the type of norm, C is the triggering condition, A and E show the activation and expiration period of the norm, and S and R indicate a reward or sanction. For example, this is an example of a possible norm: $(\langle \text{prohibition, smoking, } -, -, \text{payfine, } - \rangle, \delta)$, which is always valid since there is no duration on activation, A , and expiration, E .

All of possible norms are initialized at the beginning of the simulation with the certainty value of zero. Agents update their norms by increasing or decreasing the certainty value of each norm after making an observation. For instance, if the agent receives a fine after smoking, it will update its current value of (δ) in the above norm example with $(\delta + \epsilon)$, where ϵ is a user defined value.

An agent's current norms are used to update its beliefs and desires. The updating procedure is shown in Equations 27.3 to 27.5. Here, norms are abbreviated as N instead

of $\langle \Delta, C, A, E, S, R \rangle$. If there are any relevant rewards R (or sanctions S), the positive desire D^+ (or a negative desire D^-) will be updated. f functions are user defined functions.

$$((N_i, \delta_N), (B[\alpha]\varphi, \delta_\phi)) \Rightarrow (B[\alpha]\varphi, f(\delta_N, \delta_\phi)) \quad (27.3)$$

$$((N_i, \delta_N), (D^+\varphi, \delta_\phi), R \neq \emptyset) \Rightarrow (D^+\varphi, f(\delta_N, \delta_\phi)) \quad (27.4)$$

$$((N_i, \delta_N), (D^-\varphi, \delta_\phi), S \neq \emptyset) \Rightarrow (D^-\varphi, f(\delta_N, \delta_\phi)) \quad (27.5)$$

As an example, if the norm $(\langle \text{prohibition, smoking, } -, -, \text{payfine, } - \rangle, 0.75)$ exists and a negative desire toward paying fine $(D^- \text{payfine}, 0.55)$, assuming the agent has just paid a fine for smoking $(S \neq \emptyset)$ with $f = \min(\max(0.75, 0.55), 1)$, the resulting updated desire would be $(D^- \text{payfine}, 0.75)$.

Figure 27.3 shows the pseudocode describing an agent's behavior for one time-step in the CSL implementation. The certainty value of beliefs and desires are initialized uniformly at random at the beginning of the scenario.

```

init(blfs, des, pln, q-tbl)
repeat
  generateIntention(blfs, des, pln) ▷ Equation 2
  updatePMatrix(maxIntention)
  if (converged-Qtbl) then
    playGame(pMatrix, neighbors)
    performAction()
    update-qTable(rew, san)
  else
    performAction()
  end if
  update-norms(rew, san)
  update-beliefs(rew, san, norms) ▷ Equation 3
  update-desires(rew, san, norms) ▷ Equation 1, 4 & 5
until agent not selected

```

Figure 27.3: CSL pseudocode
(blf=Beliefs, des=Desires, pln=Plans, rew=Rewards, san=Sanctions)

Smoking Model

Our smoking model considers three sets of factors that are known to affect human smokers: personal, social, and environmental influences. Considering the complex and challenging nature of modeling smoking behaviors, specifically the addictive property of smoking, we tried to have an inclusive model that contains as many factors as possible.

Personal

Our model includes a set of personal values which are specific to each person, and depend on their personality; Dechesne et al. (2013) use a similar set of values within their model of cultural differences that affect smoking behavior. According to the sociological theory of cultural value orientation introduced by Schwartz (2006), three types of values determine cultural differences in societies. These values are defined by three bipolar cultural dimensions that can be used to describe possible resolutions to problems confronting societies. In our model, we adopted two of these values since the third dimension is specifically for cultural differences which are negligible for our relatively homogeneous undergrad population. The two adopted values are described below:

- **Embeddedness vs. autonomy:** This determines how much an individual's preferences, feelings, and ideas are affected by others through various relationships vs. being cultivated internally.
- **Mastery vs. harmony:** This refers to the dichotomy of being ambitious, daring, and self-assertive vs. being consistent, understanding, and appreciative of the environment.

The first item is referred as **individualism** (ind), and the second one as **achievement** (ach). The third item which is not included in our model is equality. In addition to these two personal values drawn from Schwartz's sociological (or anthropological) model, three other personal values are included:

- **Regret** (rgt) - In our scenario, this value shows how much the individual is regretful about smoking and is used to model the phenomenon of addiction. The role of regret in smoking behaviors is described in Conner et al. (2006); it is related to their willingness to quit smoking or decrease their tobacco usage.
- **Health** (hlt) - As the name implies, this value shows the extent to which a person is health-conscious and pays attention to medical recommendations.

- **Hedonism** (hdn) - The pleasure-seeking aspect of one's personality. Health and hedonism were also used in a related study on smoking bans in European Union countries (Dechesne et al., 2013).

Social

The second aspect of our model is used to quantify the effects of the community on the individual. To do this, we create a synthetic friendship network for our simulated community using the method described in Wang et al. (2011) for creating human networks that follow a power law degree distribution and possess homophily, a greater number of link connections between similar nodes. The network generator uses link density (*ld*) and homophily (*dh*) to govern network formation. For our smoking model, three elements are defined to determine the homophily of a node: age, gender and undergraduate major. The nodes of the graph represent the individuals (agents) in the simulation.

Environmental

The third category of factors that affect people's smoking behavior is what they observe or encounter in their surroundings. Four items are considered in this category: others, signs+butts, advertisements, miscellaneous.

Others (oth) - One major factor that affects norm compliance is observing other people's behavior. Seeing other smokers can affect the agents' decisions to obey policies, particularly when complying with smoking cessation rules. Similar behaviors in humans have been shown to exist and are usually referred to as *observational learning*. Various studies have shown the effect of observation on smoking behaviors (e.g., Akers, Lee (1996)).

Signs + butts (sbt) - This item is specifically related to the effect of installed *No Smoking* signs, that advise people to refrain from smoking. A key research challenge here is to simulate the behavior of people in response to this type of notification. A recent study by Schultz et al. (2013) on littering in public locations shows that people tend to obey installed signs when there is no trash around the sign, but when litter exists in the vicinity, the rate of people who do not follow the signs increases significantly. Using a similar approach, we consider signs and cigarette butts together and model the influence of observed cigarette butts on a person's on-campus smoking behavior.

Advertisements (adv) - Physical advertisements can also influence smoking behaviors. These advertisements are a major part of the campus smoke-free program. This category refers to tents, fliers, billboards, catalogs, posters and banners installed permanently in different locations of campus.

Miscellaneous (msc) - This category encompasses all of the other factors that might influence a smoker's decisions. One major aspect of this category is non-physical influences, especially digital, educational, and promotional activities. Also included in this category is the role of different cessation facilities available on campus, such as workshops and nicotine replacement therapy (NRT).

The five elements introduced for the personal values, the social element, and the four environmental factor are all defined as ranging from 0 to 100. The main smoking-value (SV) is calculated using this formula:

$$SV = (k_1 * ind' + k_2 * ach' + k_3 * rgt + k_4 * hlt' + k_5 * hdn + k_6 * frd + k_7 * oth + k_8 * sbt + k_9 * adv + k_{10} * msc) / \sum_{i=1}^{10} k_i \quad (27.6)$$

The smoking-value (SV) falls between 0 to 100. In this formula, k_1 to k_{10} show nine coefficients that are assigned to the user. Prime (') means complement, which in this case is equal to: "100 -". The friendship value (frd) is determined using the social model.

Agent-based Model

The original version of the agent-based model (ABM) used in this work was built to study the transportation patterns of people and vehicles (Beheshti, Sukthankar, 2012, 2014a). Before presenting the new components, we will first describe the function of the base ABM. The model was built to simulate the movement patterns of students at the University of Central Florida. The data for building the model were gathered through an online survey. In the survey, participants were asked to answer questions about the time they arrive and depart campus, locations they visit, and frequency of their visits. A set of statistical distributions was fit to the answers of each question. These distributions were then used to initialize the model parameters, including those that govern the activities of an agent.

Each agent arrives, visits locations on campus, and then leaves campus according to its own personal schedule. Various specialized rules were added to the model to improve the verisimilitude of the whole system. Examples of defined rules include limitations on the number of cars that can enter a parking lot or the hours that shuttle services operate. The accuracy of the ABM was measured in several different ways, including comparing the obtained statistics from the ABM with other independently collected data sources.

To implement the smoking simulation scenario, the proposed smoking model was added to the original ABM. We added two parameters, age and gender, to each agent's

parameter set to be used for measuring homophily in the social model. Each agent is initialized as a smoker or non-smoker at the start of the ABM, based on the number of smokers in the survey data. The smoke-free campus policy is assumed to be in effect immediately after the start of the simulation.

Having a detailed transportation model facilitates implementing the environmental aspects of the proposed smoking model in high fidelity. The assumption is that each smoker agent smokes an average of 15 (for men) and 10 (for women) cigarettes per day. These numbers are based on the reported statistics in Burns et al. (2003). The effect of observing others smoking on campus is incrementally aggregated for each agent through the described reinforcement learning algorithm. The observation occurs whenever an agent is close to an agent that is smoking at the same time.

The exact location of no-smoking signs and physical advertisements are defined in the campus map used in the ABM. Based on our observational study of the campus, cigarette butt locations are marked near the large college buildings, but not general buildings like the student union and library. This trend might occur because of the frequent cleaning of these areas, or the tendency of people to avoid smoking in heavily crowded areas. While the agent moves around campus, it passes physical advertisements. Similar to observing others smoking, every encounter with an advertisement increases its chance of affecting the user.

Figure 27.4 shows the user interface of the agent-based model. In this figure, the location of buildings, routes and also the advertisements can be seen. The last item of the environmental model (misc factors) is implemented by a random value that represents the aggregation of all other factors.

LNA setup

Here, we describe the details of implementing smoking factors on the agent-based model that follows the LNA architecture.

Personal - Personal values were added to the set of parameters possessed by each agent in the ABM. These values are calculated using distributions fitted to the available survey data.

Social - To implement the diffusion of smoking behaviors in the friendship network, a game-theoretic approach (Easley, Kleinberg, 2010) is used. Here, a simple two by two matrix is defined that contains four different states that can occur in the smoking scenario. Table 27.2 shows this matrix. The descriptions below the table show how the payoffs are calculated. The abbreviations on the right side of the equations relate to being a smoker (s) or non-smoker (n).

Each individual is either a smoker or non-smoker. The payoff for each of four



Figure 27.4: Screenshot of the agent-based model. The advertisements (orange pentagons) and no-smoking signs (red triangles) are shown on the map.

entries of a node is calculated according to three factors: personal values, network neighbors, and whether the subsequent state is similar to the current state. Similar to the mechanism that was described earlier, α and β values are added to the model to show the tendency of people to maintain their current state. These two parameters are constant positive values which make the value of the payoff higher for the cases that the agent remains a smoker or non-smoker than in the cases that a state transition occurs. The final value for the friendship element of model (frd) is calculated based on the current state of the individual and her friends, using the payoff matrix.

Environmental - Each of the four elements is represented in the model with values ranging from 0 to 100. These factors are modeled as part of agents' beliefs for the CSL architecture. For the implementation of environmental factors on LNA architecture, a simplified version of Q-learning is used to govern the effects of the environmental factors. As Table 27.3 shows, when encountering an environmental factor such as a banner, the state of an agent is defined by the current value of its personal and social elements. The agent can either be affected by the environmental factor or disregard it. In case of the first action, the value of that environmental factor will increase by a fixed amount, but in the second case nothing changes. The reward that agent receives from

		Node B		
		Smoker	Non-smoker	
Node A	Smoker	ss+ α	sn	ss = ind' + ach' + hlt' + hdn'
	Non-smoker	ns	nn+ β	sn = ind + ach + hlt + hdn'
		ss+ α	ns	ns = ind + ach + hlt' + hdn
		sn	nn+ β	nn = ind' + ach' + hlt + hdn'

Table 27.2: Payoff matrix governing the diffusion process in the friendship network. Prime (') means complement, which in this case is equal to: "100 -". ind: individualism; ach: achievement; hlt: health; hdn: hedonism.

States	current value of personal and social elements
Actions	pay attention or not
Rewards	calculated based on the values of regret, health and hedonism

Table 27.3: Q-learning definitions for state, actions, and rewards. If the agent does not pay attention, it means that the agent opts to ignore a specific environmental element. Regret and health affects the reward value positively, and hedonism affects it negatively.

each action is calculated based on three elements of its personal value vector: regret, health and hedonism. The reward value falls between -1 and +1, and is calculated using the following formula:

$$\text{reward} = (\text{regret} + \text{health} - 2 * \text{hedonism})/200 \quad (27.7)$$

A dynamic learning schedule is utilized for the Q-learning, which results in a higher rate of learning at the beginning of the simulation, and a lower one afterwards.

CSL setup

Here we describe the elements of the CSL architecture.

Beliefs, Desires, and Intentions - The two first personal values, individualism and achievement, are implemented as fixed value elements of beliefs (Type 1). The remaining three personal factors, regret, health and hedonism, plus environmental factors are implemented as variables, and part of each agent's beliefs. The certainty values (δ) for beliefs and desires are assigned uniformly at random at the beginning of the scenario. The intentions are determined according to Equation 27.2. The main desires and intentions defined in this system refer to smoking and not smoking.

Payoff Matrices - An agent plays games with both its friends and other agents near to it to determine its actions. For each action, an agent has a two by two payoff

matrix that determines the agent's decision. The agent picks the intention with the highest certainty value. The values of this payoff matrix are determined by the certainty degree of the selected intention. This means that in our architecture, the intentions do not directly determine agent's actions, instead they define payoff matrix values. The friendship value (frd) in the smoking model is calculated using the payoff matrix values.

Norm Recognition - The learning component is implemented using the Q-learning algorithm. Actions are the action performed by the agent: to smoke or refrain. The reward value is assumed to be the same as the reward value defined for the reinforcement learning and smoking diffusion in LNA. The current values of the payoff matrices determine the states of the Q-table. The selected action modifies the certainty value of norms. After an agent performs an action, it observes the consequences of its action to compute the overall received payoff, which is then used to update the Q-table.

Norms - Norms are created using the same procedure introduced. Only dynamic (variable) parts of beliefs are updated. All possible norms are initialized as having a certainty value of zero. During initialization, we create all of possible norm combinations based on the introduced norm representation: $\langle \Delta, C, A, E, S, R \rangle$. The type of norm and its reward or sanction nature can be determined by the value for C . We assume that all norms are always valid during the experiment, so we don't need to take A and E into account. Thus 12 possible norms are defined for this scenario: |obligation, prohibition, permission|*|smoking, not smoking|*|reward, sanction|.

Data

Our agent-based model uses data from three surveys of UCF students. In Spring 2012, we did an online survey of 1003 students to collect the data used to model campus transportation patterns. The other two surveys were conducted by Health Services; one of them was done in Fall 2011, before the smoke-free policy was instituted, and the second in Fall 2012, at the end of the first year of the smoke-free campus. Both of these surveys were performed as part of the annual university ACHA-NCHA reporting process. The student answers to five questions in the first survey were used to determine the numerical values for the five personal values. The personal values and corresponding survey questions are:

- **Individualism** - Do you think breathing smoke-free air on campus is a right?
- **Hedonism** - Do you think smokers have the right to smoke on campus?
- **Achievement** - Would you feel comfortable asking someone to put out their cigarette?

- **Health** - Would a smoke-free campus policy make campus healthier?
- **Regret** - If you smoke, are you interested in attending a smoking cessation program?

The questionnaire was designed using a Likert scale. The personal values in our work were matched to questions after the survey was conducted, and normal distributions fitted to the data were used to initialize the agents' personal values in the ABM. The university administration used the answers to the following three questions to determine the success of the smoke-free campus policy. In our work, the answers to the second and last question were used to show the accuracy of the proposed model. These three questions are:

- Do you support the campus smoke-free policy?
- Do you smoke?
- Are you likely to take smoking cessation classes?

The other data used to implement the model, including the location of advertisements and installed no-smoking signs, were obtained from campus sources.

Experiments

Validation is a major challenge while evaluating ABMs—how to show that the model matches reality. One approach is to evaluate the model by comparing the statistics obtained from the model with other sources of data as indicators of ground truth. Here, the data obtained from the second and third questions of the survey described in the previous section are used to evaluate the model. These two questions show the percentage of smokers among the students, and also the percentage of those who are willing to attend smoke cessation workshops.

The ABM is initialized with the same number of smokers and people willing to participate in smoking cessation classes as indicated in the survey data.¹ According to our definition, a smoker is an agent whose smoking-value, (SV), is below the quitting threshold. Similarly, we use the middle part of the proposed smoking-value range to identify an agent who is willing to attend smoking classes. An agent who is willing to participate in classes has a smoking-value between the two proposed thresholds. The assumption is that the adoption phase in the proposed architecture shows the situation where the agent has not reached the compliance phase. So, assuming that an agent in

¹Since the total number of students is known, the percentage values also determine the numbers, hence we use the terms interchangeably.

the compliance mode is willing to attend smoking classes is consistent with the proposed architecture, because attending class is not a clear quitting task, but is a behavior toward quitting (the action phase).

Table 27.4 shows the parameters that are used in the experiments to determine the smoking range. As the table shows, the value 50 is used for the first threshold and 90 for the second threshold shown in Figure 27.1. In our experiments, the values for the coefficients k_3 , k_4 and k_6 in equation 27.6 were 3, 3 and 2. The other coefficients were equal to 1. In the next section, a set sensitivity analysis experiments related to these values are presented. For the network generation part, the values for the link density, ld , and homophily, dh , were 0.40 and 0.66.

Agent State	Range
Non-smoker	90–100
Willing to participate in classes	50–90

Table 27.4: Experimental settings for smoking-value (sv)

In addition to the LNA and CSL architectures, we have also implemented the NBDI architecture (Norm Belief-Desire Intention) (Criado et al., 2010a). The NBDI benchmark does not play the social dilemma game and does not use reinforcement learning to generate and update norms. In this case, intentions determine actions, and then the norms are updated based on the feedback received from the environment. Note that the way that the norm representation was implemented (by modifying the certainty value of norms) is not part of the original version of NBDI. The norm recognition part in the original NBDI was assumed to work as a blackbox, and there was insufficient detail about its implementation to recreate it. Hence we simply used the same norm recognition structure for both CSL and NBDI.

Results

Using these assumptions, we ran our agent-based simulation for a period of a year from Fall 2011 to Fall 2013. In these experiments, we initialized the simulation with the same number of smokers and students willing to go to the classes as the initial survey data, and then compared the numbers obtained from the simulation with the final survey data. During this period, the agents commute to campus and follow schedules governed by the transportation model. The proposed smoking model simulates the smoking behavior of students during the year of study. The average simulation error of ten runs of the model are reported in Figure 27.5. Simulation error refers to the difference between the values obtained from each method and the real value from the experimental data. The two measures shown here are the percentage of smoker students

and the percentage of smoker students who are willing to attend smoking cessation classes. The empirical data for the percentage of smokers was also available for 2013.

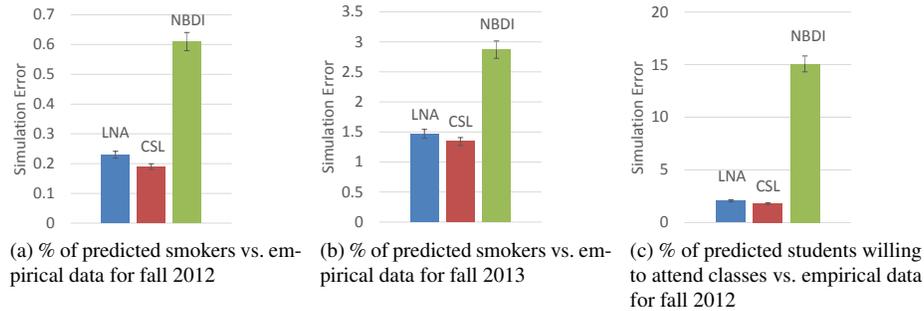


Figure 27.5: Comparison between the performance of different normative architectures. The simulation error refers to the difference between the value obtained by each method and the empirical survey data.

Figure 27.5 shows the comparison between the number of students who were smokers and students willing to participate in smoking cessation classes. The performance of CSL at predicting the actual adoption of the smoking cessation norm is comparable to the specialized smoking model (LNA) and superior to NBDI.

A powerful feature of agent-based models is their potential ability of predicting future trends, which is not possible if past data is not representative of future data. This can be a great tool for policy makers who want to analyze the effects of modifying various parameters of a specific model. In Figure 27.6, the predicted percentage of smokers for the period of the years 2011 to 2016 is shown. The values shown for the years 2011 to 2013 are the same as Figure 27.5. Both models slightly overpredict the number of smokers identified using the UCF Health Services survey data, with CSL being slightly more accurate. For the year 2013, it was confirmed by the Health Services department that the reported rate (3.9 %) seems a bit lower than what they were expecting based on national and state averages. One possibility is that smoking behavior is underreported by the students or being supplanted by vaping. The current assumption in our model is that different system properties remain the same during the simulated years. A factor that our model does not take into account is the gradual change of the population as students arrive to the school and graduate.

Table 27.5 shows a comparison between the different architectures at predicting the perceived social unacceptability of smoking. This phenomenon is reported in many smoking studies including Dotinga et al. (2005) and Hammond et al. (2006) as occurring when smoking bans exist in human cities. Brown et al. (2005) showed that perceived social acceptability of smoking among referent groups is independently as-

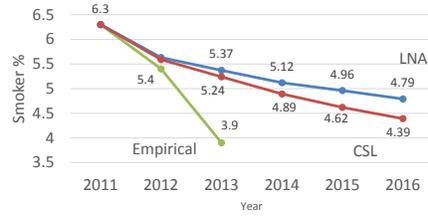


Figure 27.6: Predicted percentage of smokers for future years

sociated with both strength of intention to quit and actual quitting behavior.

	Beta	p level
CSL	0.22	0.001
LNA	0.001	0.007
NBDI	-0.01	0.005

Table 27.5: Standard coefficient (Beta) values of the applied linear regression to perceived social acceptability of smoking (independent variable) and quitting intention (dependent variable)

In our smoking model, it is assumed that an agent has the intention to quit smoking if its smoking value (SV) is within the first and second threshold values. The social unacceptability of smoking across the population of agents is determined using the value for one of the agent’s personal characteristics (IND). The value of this factor was initialized based on data from a survey question asking whether the participant believes smoking is acceptable on campus. A linear regression model was used to examine the relationship between these two elements, and the standard coefficient (Beta) value of the applied linear regression is shown in Table 27.5. The CSL model produces a positive Beta value, which is consistent with the real-world data. This shows that, using CSL, agents are able to reason about the socially perceived unacceptability of smoking behavior, and modify their behaviors accordingly. Therefore, CSL is modeling norm emergence in a more realistic manner. On the other hand, the Beta values for the LNA and NBDI architectures are close to zero, which does not accurately reflect the results reported in independent smoking studies.

Additionally, we performed a set of sensitivity analysis experiments on the results that we obtained from the two architectures. Since our models include a number of variables that could directly affect the final behavior of our system, the sensitivity analysis illuminates the effect that each variable can have on the final outcome. Five of the ten coefficients (Equation 27.6) plus the two threshold values for determining the three

stages of norm formation (Figure 27.1) are used as the independent variables in our sensitivity analysis model. The remaining five coefficients are not shown due to their close relation to the current coefficients. The analysis is performed on one independent value at a time.

Figure 27.7 shows the range of output values for different values that can be assigned to five of the k_i coefficients, and similarly Figure 27.8 shows the output range for the two threshold values. By comparing the results shown in Figure 27.7a with 27.8a, and also 27.8a with 27.8b, we can observe that LNA seems to be more sensitive to noise than CSL. By changing the coefficient values from 0 to 6, the maximum change in the percentage of smokers is close to 4 for LNA, and less than 2 for CSL. In case of the two threshold values (shown in Figure 27.8) LNA's results varies in the range of 3.5, while the length of CSL's range is less than 2.5. Overall, the sensitivity of the model's output to the set of input values is low and because of type of equation that we proposed, the output range for different values remains linear.

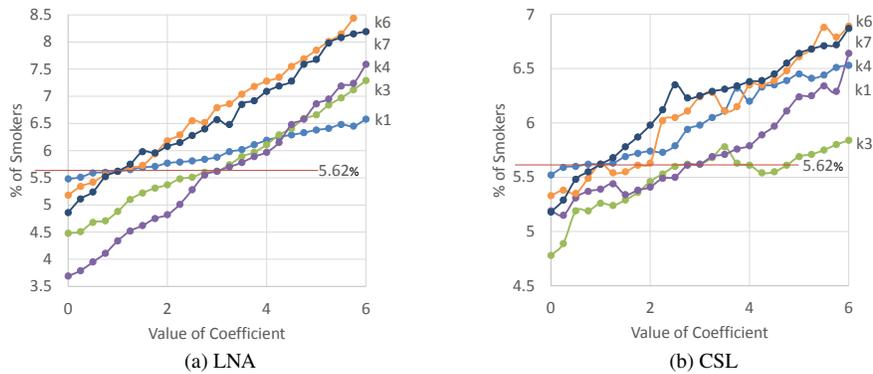


Figure 27.7: Sensitivity analysis of the values for five coefficient values in the regression equation introduced for determining the final smoking value in our models. Horizontal red lines show the current values reported by our models.

We also study ablated versions of the model that lack one of the three smoking elements (social, environmental, or personal). The results for alternate months during the year of simulation are reported in Figure 27.9. The reported results are, again, averaged over ten runs, and in all cases the initialization configuration is based on the survey data. In Figures 27.9a and 27.9b the left red star shows the starting value which is the empirically measured value, and is the same for all four experiments. Without the personal and environmental components, the model tends to underestimate results in comparison to the final empirical results. Without the social part, the model overestimates smoking behavior. Based on the size of differences between the empirical

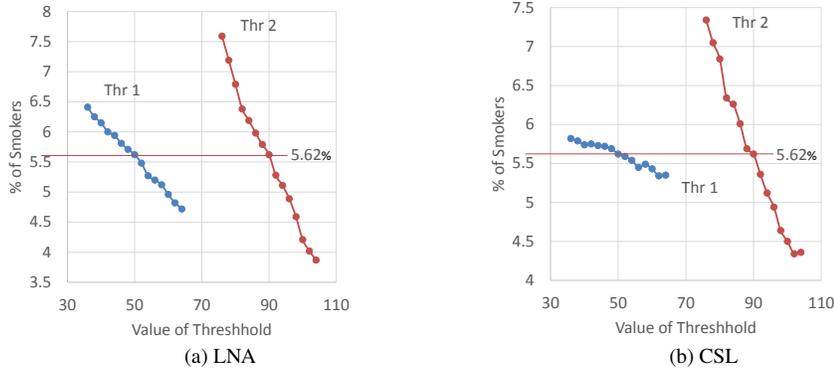


Figure 27.8: Sensitivity analysis of the effects of the two threshold values in our models. Horizontal red lines show the current values reported by our models.

results and the other experiments for CSL, it can be concluded that the personal values are the major predictors in determining smoking behaviors. Environmental factors had the lowest impact on predicting smoking behavior.

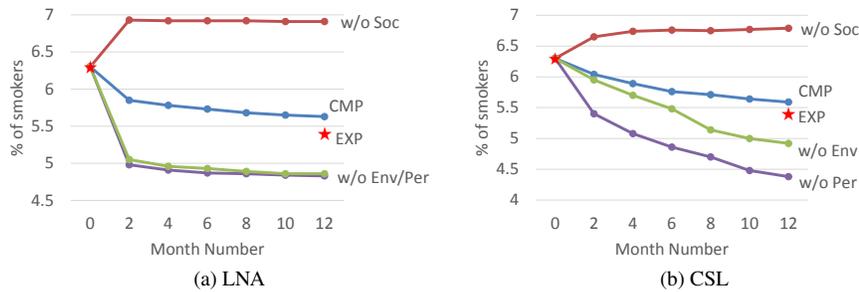


Figure 27.9: The percentage of smoker students in LNA (a), and in CSL (b) during the one year simulation period. The numbers from the survey data are marked by the red star icons at the beginning and end of the simulation period (Experimental/EXP). The figure shows the predictions of the proposed model (complete/CMP), the model without the personal values, without the social aspect, and without environmental influence. There is a close match between the predicted values of the complete model and the survey data.

Discussion

The LNA architecture is based on a simple model of norm adoption that is similar to many of the mechanisms that are currently used for constructing normative multi-agent

systems. On the other hand, CSL represents a group of architectures that employ more complex structures for normative reasoning. CSL integrates internal cognitive structures with social interaction mechanisms. Our results demonstrate that CSL provides a richer model of health-related social norms. Specifically, when it comes to modeling the intricacies of humans' behaviors – like the correlation between the unacceptability of smoking in a society and quitting intention – our lightweight normative architecture exhibits a greater sensitivity to design assumptions and produces a lower fidelity model, than CSL.

Conclusion

In this chapter, we introduced two architectures (LNA and CSL) for modeling the emergence of norms in agent-based models. Our proposed normative architectures focus on the impact of social factors on smoking behavior and are intended to complement biological models of the effects of nicotine addiction. We examine the effect of a richer reasoning model on the overall performance of a model in understanding human behavior. Although both methods demonstrate good prediction abilities at the population level, the agents with more advanced cognitive and learning structures exhibit a higher level of verisimilitude in simulating human behavior. This was a gratifying result because we expect agent-based modeling to have this potential. At the same time, results are driven by estimates of a number of parameter values, and more data is needed to permit developing better methods for making such estimates and then updating them as additional information comes in. Doing so would be akin to what statisticians do, but we would anticipate results to be much better with the cognitive and learning structures of the ABM.

Author Biography

Gita Sukthankar is an Associate Professor with the Department of Computer Science at University of Central Florida and an affiliate faculty member at UCFs Institute for Simulation and Training. She directs the Intelligent Agents Lab (<http://ial.eecs.ucf.edu/>). Rahmatollah Beheshti is a Postdoctoral Fellow with the School of Public Health, Johns Hopkins University and a former member of the Intelligent Agents Lab.

Acknowledgments

This research was supported by NSF IIS-08451 and DARPA HR001117S0018-SocialSim-FP-026.

Bibliography

- Akers Ronald L, Lee Gang.* A Longitudinal Test of Social Learning Theory: Adolescent Smoking // *Journal of Drug Issues.* 1996. 26. 317–343.
- Beckman Richard, Kuhlman Christopher, Marathe Achla, Nsoesie Elaine, Swarup Samarth.* Modeling the Spread of Smoking in Adolescent Social Networks // *Proceedings of the Fall Research Conference of the Association for Public Policy Analysis and Management.* Washington DC, 2011.
- Beheshti Rahmatollah, Ali Awrad Mohammed, Sukthankar Gita.* Cognitive Social Learners: an Architecture for Modeling Normative Behavior // *Proceedings of the AAAI Conference on Artificial Intelligence.* Austin, TX, Jan 2015.
- Beheshti Rahmatollah, Sukthankar Gita.* Extracting Agent-Based Models of Human Transportation Patterns // *Proceedings of the ASE/IEEE International Conference on Social Informatics.* Washington, D.C., dec 2012. 157–164.
- Beheshti Rahmatollah, Sukthankar Gita.* A Hybrid Modeling Approach for Parking and Traffic Prediction in Urban Simulations // *AI and Society: Journal of Knowledge, Culture and Communication.* 2014a.
- Beheshti Rahmatollah, Sukthankar Gita.* A Normative Agent-based Model for Predicting Smoking Cessation // *Proceedings of the International Conference on Autonomous Agents and Multi-agent Systems.* Paris, France, may 2014b. 557–564.
- Brown D.G., Riolo R., Robinson D.T., North M., Rand W.* Spatial process and data models: Toward integration of agent-based models and GIS // *Journal of Geographical Systems.* 2005. 7, 1. 25–47.
- Burns David M, Major Jacqueline M, Shanks Thomas G.* Changes in number of cigarettes smoked per day: Crosssectional and birth cohort analyses using NHIS // *Smoking and Tobacco Control Monograph.* 2003. 15. 83–99. NIH publication no. 03-5370.

- Casali Ana, Godo Lluis, Sierra Carles.* A Logical Framework to Represent and Reason about Graded Preferences and Intentions. // Eleventh International Conference on Principles of Knowledge Representation and Reasoning. 2008. 27–37.
- Conner Mark, Sandberg Tracy, McMillan Brian, Higgins Andrea.* Role of anticipated regret, intentions and intention stability in adolescent smoking initiation // British Journal of Health Psychology. 2006. 11, 1. 85–101.
- Criado N., Argente E., Noriega P., Botti V.* Human-inspired model for norm compliance decision making // Information Sciences. 2013. 245, 0. 218 – 239. Statistics with Imperfect Data.
- Criado Natalia, Argente Estefania, Botti V.* A BDI architecture for normative decision making // Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems. 2010a. 1383–1384.
- Criado Natalia, Argente Estefania, Botti Vicent.* Normative deliberation in graded BDI agents // Multiagent System Technologies. 2010b. 52–63.
- Dacunto Philip J, Cheng Kai-Chung, Acevedo-Bolton Viviana, Klepeis Neil E, Repace James L, Ott Wayne R, Hildemann Lynn M.* Identifying and quantifying secondhand smoke in multiunit homes with tobacco smoke odor complaints // Atmospheric Environment. 2013.
- Dechesne Francien, Di Tosto Gennaro, Dignum Virginia, Dignum Frank.* No smoking here: values, norms and culture in multi-agent systems // Artificial Intelligence and Law. 2013. 21, 1. 79–107.
- Dotinga Aafje, Schrijvers Carola TM, Voorham Anthonius JJ, Mackenbach Johan P.* Correlates of stages of change of smoking among inhabitants of deprived neighbourhoods // The European Journal of Public Health. 2005. 15, 2. 152–159.
- Easley David, Kleinberg Jon.* Networks, crowds, and markets. 8. 2010.
- Gutkin Boris S., Dehaene Stanislas, Changeux Jean-Pierre.* A neurocomputational hypothesis for nicotine addiction // Proceedings of the National Academy of Sciences. 2006. 103, 4. 1106–1111.
- Hammond David, Fong Geoffrey T, Zanna Mark P, Thrasher James F, Borland Ron.* Tobacco denormalization and industry beliefs among smokers from four countries // American Journal of Preventive Medicine. 2006. 31, 3. 225–232.
- Homer Jack B, Hirsch Gary B.* System dynamics modeling for public health: background and opportunities // American journal of public health. 2006. 96, 3. 452–458.

- Levy David T, Nikolayev Leonid, Mumford Elizabeth.* Recent trends in smoking and the role of public policies: results from the SimSmoke tobacco control policy simulation model // *Addiction*. 2005. 100, 10. 1526–1536.
- Luo Sean X, Covey Lirio S, Hu Mei-Chen, Levin Frances R, Nunes Edward V, Winhusen Theresa M.* Toward personalized smoking-cessation treatment: Using a predictive modeling approach to guide decisions regarding stimulant medication treatment of attention-deficit/hyperactivity disorder (ADHD) in smokers // *The American Journal on Addictions*. 2015.
- Schultz P Wesley, Bator Renée J, Large Lori Brown, Bruni Coral M, Tabanico Jennifer J.* Littering in Context Personal and Environmental Predictors of Littering Behavior // *Environment and Behavior*. 2013. 45, 1. 35–59.
- Schwartz Shalom H.* A theory of cultural value orientations: Explication and applications // *International Studies in Sociology and Social Anthropology*. 2006. 104. 33.
- Sen Sandip, Airiau Stéphane.* Emergence of norms through social learning // *Proceedings of the International Joint Conference on Artificial Intelligence*. 2007. 1507–1512.
- Timms Kevin P, Rivera Daniel E, Collins Linda M, Piper Megan E.* System identification modeling of a smoking cessation intervention // *16th IFAC Symposium on System Identification, Brussels, Belgium*. 2012. 11–13.
- Wang Xi, Maghami Mahsa, Sukthankar Gita.* Leveraging network properties for trust evaluation in multi-agent systems // *Proceedings of the IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology-Volume 02*. 2011. 288–295.