

A Normative Agent-based Model for Predicting Smoking Cessation Trends

Rahmatollah Beheshti
Department of EECS
University of Central Florida
Orlando, FL, USA
beheshti@knights.ucf.edu

Gita Sukthankar
Department of EECS
University of Central Florida
Orlando, FL, USA
gitars@eecs.ucf.edu

ABSTRACT

Norms are an important part of human social systems, governing many aspects of group decision-making. Yet many popularly used social models neglect to model normative effects on human behavior, relying on simple probabilistic and majority voting models of influence diffusion. Within the multi-agent research community, the study of norm emergence, compliance, and adoption has resulted in new architectures and standards for normative agents; however few of these models have been successfully applied to real-world public policy problems.

In this paper, we propose a new lightweight architecture for constructing normative agents to model human social systems; the aim of our research is to be able to study the effects of different public policy decisions on a community. Here we present a case study showing the usage of our architecture for predicting trends in smoking cessation resulting from a smoke-free campus initiative. Our agent-based model combines social, environmental, and personal factors to accurately predict smoking trends and attitudes. The performance of both the whole and ablated model is evaluated against statistics from an independent source.

Categories and Subject Descriptors

I.2.11 [Distributed artificial intelligence]: Multi-agent systems

Keywords

norms; agent architectures; agent-based modeling; smoking cessation

1. INTRODUCTION

One barrier to creating realistic large-scale models of human social systems is the lack of good general purpose computational models of human interactions; without such models, it is impossible to accurately account for the intricate action dependencies engendered by both explicit and implicit interpersonal communications. However research on special purpose human interaction models has flourished, bringing a greater understanding of the computational processes underlying teamwork [32], information diffusion [19],

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*
Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

and adversarial situations [11]. Armed with these tools, social scientists have been able to mathematically describe more complicated social phenomena. Similarly, we believe that the research on computational models of norms and normative agent architectures is ripe for greater inclusion in social simulations.

This paper describes a lightweight architecture for simulating normative effects using agent-based models. The overarching aim of our research is to create a general purpose agent-based modeling (ABM) and simulation system for studying the effects of public policy decisions on a large range of social phenomena, including personal health decisions, sustainability behaviors, and opinion formation. Norms are an important key to understanding the function of human groups, teams, and communities; they are a ubiquitous but invisible force governing many human behaviors. Bicchieri describes human norms as: “the language a society speaks, the embodiments of its values and collective desires, the secure guide in the uncertain lands we all traverse, the common practices that hold human groups together.” [8]

A *normative agent* refers to an autonomous agent who demonstrates normative behavior; these agents must be able to reason about the norms with which they should comply, and occasionally violate them if they are in conflict with each other or with the agent’s private goals [22]. For individual agents, reasoning about social norms can easily be supported within many agent architectures; Dignum [16] defines three layers of norms (private, contract, and convention) that can be used to model norms within a BDI framework. At the population level, norm emergence, whether a group of agents converges to a consistent set of norms, is an interesting question, and both theoretical and computational models have been presented to describe norm emergence in social systems [31, 34]. Previous work on norms, such as the EMIL project [21], has shown promising results on modeling real-world phenomena such as traffic patterns, Wikipedia article authorship, and financial decisions. Here we seek to integrate normative effects with other types of human behavior models to produce a more comprehensive picture of human communities, rather than limiting our analysis to norms alone. Hence our proposed ABM simulates both environmental and network effects, in combination with norms.

Human social systems tend to be complex by nature; our philosophy is that constructing multi-layered models is of paramount importance when simulating real-world scenarios, since it is unlikely that a single type of interaction model will correctly account for all the observed effects. This paper presents an ablative study showing the relative contribution

of the different layers of the ABM on predicting the impact of a smoke-free campus initiative on student smoking cessation behavior. Section 2 describes the related work on normative agents and summarizes other modeling efforts on smoking cessation. Our proposed model to simulate smoking behaviors includes three factors: 1) personal values, 2) social networks, and 3) environmental influences; a detailed description is provided in Section 3. The norm in the smoking case study is the acceptability of smoking on a smoke-free campus. Agents modify their beliefs based on a combination of personal, environmental and social factors. The normative model is operationalized as part of an activity-oriented microsimulation of transportation patterns on a large university campus. Inclusion of a detailed transportation model facilitates simulating propinquity effects that arise from physical proximity. Section 4 presents results on the performance of our model at predicting smoking cessation attitudes. Although this paper focuses on smoking behavior, we believe our architecture is sufficiently general to study a variety of public policy scenarios. For instance, influencing norms through targeted climate change messaging has been shown to be a powerful tool for stimulating sustainable behaviors [28].

2. RELATED WORK

This section provides an overview of the process of creating social systems with normative agents before describing the related work on smoking cessation. Norm adoption and compliance are key to the study of normative agents. The general assumption behind norm adoption is that an agent will adopt another agent’s goal, on the condition that the adopter comes to believe that the achievement of the adoptee’s goal will increase its chances of achieving a previous held goal [3]. Castelfranchi describes two types of norm adoption: 1) instrumental, in which agents are motivated to obey a norm that benefits them and 2) terminal, which implies that the agents do not have any other choice other than following the norms [12].

Norm compliance usually refers to the process by which a normative belief becomes a normative goal [12]. Note that adoption is not synonymous with compliance in norms. An agent may adopt to a norm but choose to violate that norm later. For instance, agent transgressions can occur when the expected rewards obtained with detection surpass the expected rewards obtained by being norm-compliant [18].

The existence of norm conflicts raises the possibility of norm violations. As [4] points out, norms may be conditioned on a variety of factors including spatial, temporal, cultural and social circumstances. Norm violation is the byproduct of having a flexible norm system. In a hard-wired system in which the norms are fixed and the agents must comply, it is impossible to have violation and conflicts. Accordingly, various conflict resolution techniques have been used in the literature. Some of these methods are similar to the techniques used in general multi-agent systems, but many are specific to normative domains. For instance, a meta-norm usually refers to a higher level norm that agents consult in case of conflicts. A meta-norm can be as simple as selecting a norm at random when a conflict occurs or can be a much more complex resolution procedure. Norm conflict can be also dealt with using argumentation-based approaches (e.g.,[24],[25]).

A fundamental research question is how norms emerge in social systems. One approach is to model this phenomenon through the use of learner agents that adapt their behavior based on sanctions and rewards. Sen and Airiau’s work [31] in this area, in which agent interactions are modeled using payoff matrices, inspired much subsequent research on norm emergence through social learning in agent societies. A recent extension which adds network structure to the social system is described in [34].

Outside of computer science, the social norm marketing approach has become an important tool for public health messaging [1]. There the emphasis is on changing human social norms, rather than computationally modeling them. These types of methods have been very successful at curbing college drinking and substance abuse [26]. This indicates that our proposed approach of building normative effects into our model should be highly effective, given the previously demonstrated relevance of norms to human smoking behavior.

Non-normative models of smoking behavior already exist; for instance, *SimSmoke* is one of the most 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 [20]. Other types of simulations have been used to model the consequences of second-hand smoking [14]. In addition to norms, our proposed approach also simulates network effects as was done in Beckman et al.’s study on the propagation of adolescent smoking behavior [5]. 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 emergence of norms in different countries after the commencement of anti-smoking legislation [15]. 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.

3. MODEL

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. Our proposed model considers three factors that are known to affect human smokers: personal, social, and environmental influences, and our lightweight architecture is described in the following section.

3.1 Architecture

Our architecture encapsulates some of the functionality of earlier normative architectures while remaining simple and lightweight. One oft-cited previous work in this area, the BOID architecture, extends the classic BDI approach to include a fourth element—the notion of obligation [9]. The idea of obligation was introduced into the architecture to support social commitments, such as norms. Norms can be viewed as following a three stage life cycle, including formation, propagation, and emergence [27]. Adding norm emergence provides scalability and flexibility to normative environments. The EMIL framework [21] was introduced after the BOID architecture and represents the culmination of extensive research on norm emergence. Similar to BDI, EMIL

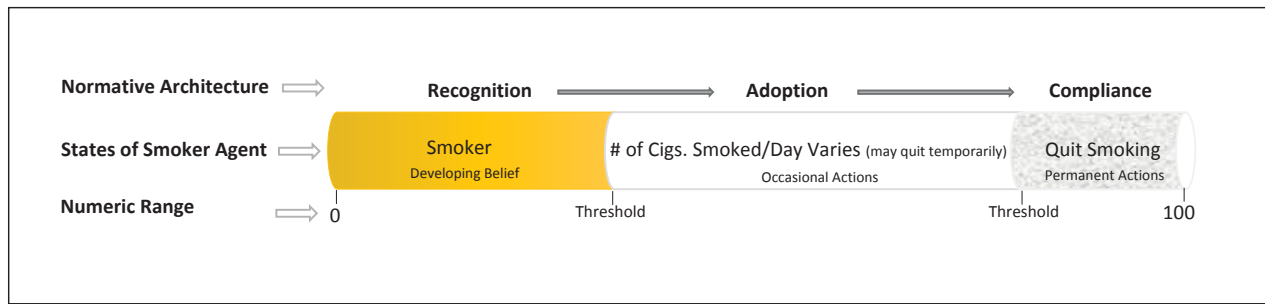


Figure 1: A schematic representation of our proposed architecture. The top row shows the three stages of the normative architecture. The middle row presents the observations corresponding to the stages within the context of the smoking scenario. The smoking norm life cycle is governed by a parameter (*smoking-value*) ranging from 0 to 100. The two user-defined thresholds (bottom row) determine 1) when an agent enters each stage and 2) what transpires.

uses belief, goal, intention and action as the procedure for norm emergence. These components are defined within two parent-categories: 1) epistemic (responsible for recognizing norms) and 2) pragmatic (responsible for behavior based on normative representation). Using the EMIL architecture in real scenarios can be challenging due to the elaborate design of its cognitive mechanisms, so we propose the following simplified architecture for how norms affect smoking behavior.

Each agent has a personal *smoking-value* ranging from 0 to 100 that governs its behavior. As shown in Figure 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. As described in the literature, 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. In order 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 situation when the agent really starts quitting. The next sections describe the factors considered by our model.

3.2 Personal

Our model includes a set of personal values which are specific to each person, and depend on their personality; Dechesne et al. use a similar set of values within their model of cultural differences that affect smoking behavior [15]. According to the sociological theory of cultural value orientation introduced by Schwartz [30], 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 [13]; 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 cares about her health, and also pays attention to medical recommendations.
- **Hedonism** (hdn) - The pleasure-seeking aspect of one’s personality. Health and hedonism were also used in the EU smoking model [15].

3.3 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 [33] for creating human networks that follow a power law degree distribution and possess homophily, a greater number of link connections between similar nodes.¹ The network generator uses link density (*ld*) and homophily (*dh*) to govern network formation. A simplified version of the pseudo-code for this method is shown in Figure 2. 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.

In order to implement the diffusion of smoking behaviors in the friendship network, a game-theoretic approach [17] is used. Here, a simple two by two matrix is defined that contains four different states that can occur in the smoking scenario. Table 1 shows this matrix. The descriptions below

¹Commonly described as “birds of a feather flock together” [23]

```

G = Null
repeat
  sample r from uniform distribution U(0, 1)
  if r ≤ ld then

    randomChooseSource(G)
    determineCandidateSink(dh, G)
    pickSink()      ▷ based on power-law distribution
    connect(source, sink)
  else
    add a new node to G
  end if
until desired number of nodes added to the network

```

Figure 2: Synthetic friendship network generator

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).

		Node B	
		Smoker	Non-smoker
Node A	Smoker	ss+ α	sn
	Non-smoker	ns	nn+ β

$$\begin{aligned}
ss &= ind' + ach' + hlt' + hdn \\
sn &= ind + ach + hlt + hdn' \\
ns &= ind + ach + hlt' + hdn \\
nn &= ind' + ach' + hlt + hdn'
\end{aligned}$$

Table 1: 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

Each individual is either a smoker or non-smoker. The payoff for each of four 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. In order to show the tendency of people to maintain their current state, α and β values are added to the model. 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.

3.4 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.

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 2: 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.

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., [2]).

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. [29] 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).

Each of these four elements is represented in the model with values ranging from 0 to 100. A simplified version of Q-learning is used to govern the effects of the environmental factors. As Table 2 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 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:

$$reward = (regret + health - 2 * hedonism) / 200 \quad (1)$$

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.

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:

$$sv = (k_1 * ind' + k_2 * ach' + k_3 * hlt + k_4 * hdn' + k_5 * rgt + k_6 * sbt + k_7 * oth + k_8 * adv + k_9 * frd) / \sum_{i=1}^9 k_i \quad (2)$$

The smoking-value (sv) falls between 0 to 100. In this formula, k_1 to k_9 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.

3.5 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 [6, 7]. 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 was gathered through an online survey. In the survey, participants were asked to answer to 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. 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 (described in the next section). We added two parameters, age and gender, to each agent's parameter set to be used for measuring homophily in the social model. (The third one, field of study, was available in the original version.) 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 [10]. 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 effectiveness.

Figure 3 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.



Figure 3: Screenshot of the agent-based model. The advertisements (orange pentagons) and no-smoking signs (red triangles) are shown on the map.

3.6 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 introduced in Section 3.2. 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 first 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, was obtained from campus sources.

4. RESULTS

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 is 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.² According to the definition presented in Section 3.1, 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 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 3 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 1. In our experiments, the values for the coefficients k_3 , k_4 and k_5 in equation 2 were 3, 3 and 2. The other coefficients were equal to 1. For the network generation part, the values for the link density, ld , and homophily, dh , were 0.40 and 0.66.

²Since the total number of students is known, the percentage values also determine the numbers, hence we use the terms interchangeably.

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

Table 3: Experimental settings for smoking-value (sv)

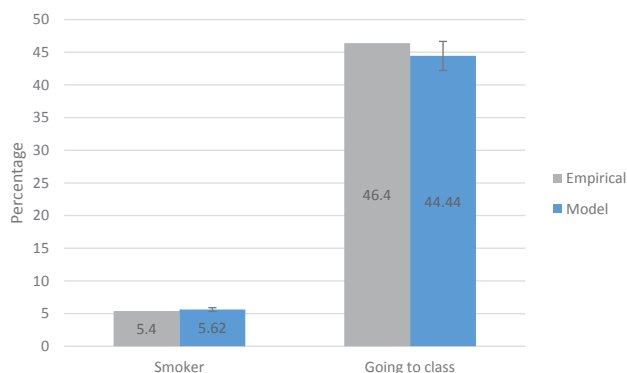


Figure 4: **Left:** the percentage of smokers in Fall 2012. **Right:** the percentage of students willing to participate in smoke cessation classes. The grey columns show the reported percentages during the survey data, and the blue ones show the percentages predicted by our model.

Using these assumptions, we ran our agent-based simulation for a period of a year from Fall 2011 to Fall 2012. 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 compare 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 results of ten runs of the model are reported in Figure 4. The figure also shows the corresponding statistics obtained from the conducted surveys. 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. As the figure shows, the model's results are very close to the reported statistics.

After evaluating the complete model, we also study ablated versions of the model that lack one of the three elements (social, environmental, or personal). The results for alternate months during the year of simulation are reported in Figure 5. The reported results are, again, averaged over ten runs, and in all cases the initialization configuration is based on the survey data. In Figures 5a and 5b 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 results and the other experiments, it can be concluded that the personal value is the major predictor in determining smoking behaviors. Environmental factors had the lowest impact on predicting smoking behavior.




Figure 5: The percentage of smoker students (a), and those who are willing to attend smoking classes (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. The figure shows the predictions of the proposed model (complete), 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.

5. CONCLUSION AND FUTURE WORK

Despite the fact that normative agent architectures have improved significantly during recent years, implementation of normative models for large, complex real-world problems has been lacking. Most existing theories and architectures have been evaluated either on artificial scenarios or on small real-world problems. In this paper, we present a lightweight normative architecture that can be initialized using survey data to model real-world scenarios and demonstrate its usage in modeling the impact of smoking cessation policies on a large university campus. We believe that our model could also be utilized (with some modifications) for similar public-policy problems in human societies.

UCF Health Services plans to promote the importance of encouraging other community members to refrain from smoking on campus. One of the measures used by the university policy makers to demonstrate the success of the smoke-free campus program was demonstrating increases in the percentage of people who feel comfortable enough to ask others to extinguish their cigarettes. Another aim is to increase the awareness of non-smoker students about the harmful effects of second-hand smoking.

Two issues remain for future work. The first is including the usage of electronic cigarettes in the model. This type of cigarette is not banned on the campus and is gaining popularity rapidly. The increasing penetration of electronic nicotine delivery systems (ENDS) and current debates among policy makers on the pros/cons of their usage have increased the importance of this topic. Our second planned extension is incorporating legal forces into our model. Smoke cessation policies usually start with encouraging approaches, and then add regulatory aspects. Similarly, this campus smoke-free policy started with a recommendation approach but is moving to a punitive one in future years. Extending the current model to predict and analyze the result of employing legal forces, including penalties and sanctions, is an interesting avenue for future work.

6. ACKNOWLEDGMENTS

The authors would like to thank the UCF's Health Services, especially Dr. Mary Schmidt-Owens, for providing sur-

vey data from the smoke-free campus program. This work is supported by the National Science Foundation, under grant IIS-08451.

7. REFERENCES

- [1] National Social Norms Institute. <http://www.socialnorms.org/index.php>.
- [2] R. L. Akers and G. Lee. A longitudinal test of social learning theory: Adolescent smoking. *Journal of Drug Issues*, 26:317–343, 1996.
- [3] G. Andrighetto and R. Conte. Cognitive dynamics of norm compliance. from norm adoption to flexible automated conformity. *Artificial Intelligence and Law*, 20(4):359–381, 2012.
- [4] G. Andrighetto, S. Cranefield, R. Conte, M. Purvis, M. Purvis, B. T. R. Savarimuthu, and D. Villatoro. (Social) norms and agent-based simulation. In S. Ossowski, editor, *Agreement Technologies*, volume 8, pages 181–189. Springer, 2013.
- [5] R. Beckman, C. Kuhlman, A. Marathe, E. Nsoesie, and S. Swarup. Modeling the spread of smoking in adolescent social networks. In *Proceedings of the Fall Research Conference of the Association for Public Policy Analysis and Management*, Washington DC, 2011.
- [6] R. Beheshti and G. Sukthankar. Extracting agent-based models of human transportation patterns. In *Proceedings of the ASE/IEEE International Conference on Social Informatics*, pages 157–164, Washington, D.C., December 2012.
- [7] R. Beheshti and G. Sukthankar. A hybrid modeling approach for parking and traffic prediction in urban simulations. *AI and Society: Journal of Knowledge, Culture and Communication*, 2014.
- [8] C. Bicchieri. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press, 2006.
- [9] J. Broersen, M. Dastani, J. Hulstijn, Z. Huang, and L. van der Torre. The BOID architecture: conflicts between beliefs, obligations, intentions and desires. In

- Proceedings of International Conference on Autonomous Agents and Multiagent Systems*, pages 9–16. ACM, 2001.
- [10] D. M. Burns, J. M. Major, and T. G. Shanks. Changes in number of cigarettes smoked per day: Crosssectional and birth cohort analyses using NHIS. *Smoking and Tobacco Control Monograph*, (15):83–99, 2003. NIH publication no. 03-5370.
- [11] C. Camerer. *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press, 2003.
- [12] C. Castelfranchi. *Cognitive and social action*. Routledge, 1995.
- [13] M. Conner, T. Sandberg, B. McMillan, and A. Higgins. Role of anticipated regret, intentions and intention stability in adolescent smoking initiation. *British Journal of Health Psychology*, 11(1):85–101, 2006.
- [14] P. J. Dacunto, K.-C. Cheng, V. Acevedo-Bolton, N. E. Klepeis, J. L. Repace, W. R. Ott, and L. M. Hildemann. Identifying and quantifying secondhand smoke in multiunit homes with tobacco smoke odor complaints. *Atmospheric Environment*, 2013.
- [15] F. Dechesne, G. Di Tosto, V. Dignum, and F. Dignum. No smoking here: values, norms and culture in multi-agent systems. *Artificial Intelligence and Law*, 21(1):79–107, 2013.
- [16] F. Dignum. Autonomous agents with norms. *Artificial Intelligence and Law*, 7(1):69–79, 1999.
- [17] D. Easley and J. Kleinberg. *Networks, crowds, and markets*, volume 8. Cambridge Univ Press, 2010.
- [18] M. S. Fagundes, H. Billhardt, and S. Ossowski. Reasoning about norm compliance with rational agents. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pages 1027–1028, 2010.
- [19] D. Kempe, J. Kleinberg, and É. Tardos. Influential nodes in a diffusion model for social networks. *Automata, Languages and Programming*, pages 1127–1138, 2005.
- [20] D. T. Levy, L. Nikolayev, and E. Mumford. Recent trends in smoking and the role of public policies: results from the SimSmoke tobacco control policy simulation model. *Addiction*, 100(10):1526–1536, 2005.
- [21] U. Lotzmann, M. Möhring, and K. G. Troitzsch. Simulating the emergence of norms in different scenarios. *Artificial Intelligence and Law*, 21(1):109–138, 2013.
- [22] M. Luck, S. Mahmoud, F. Meneguzzi, M. Kollingbaum, T. J. Norman, N. Criado, and M. S. Fagundes. Normative agents. In S. Ossowski, editor, *Agreement Technologies*, volume 8, pages 209–220. Springer, 2013.
- [23] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
- [24] S. Modgil and M. Luck. Argumentation based resolution of conflicts between desires and normative goals. *Argumentation in Multi-Agent Systems*, pages 19–36, 2009.
- [25] N. Oren, A. Rotolo, L. van der Torre, and S. Villata. Norms and argumentation. In S. Ossowski, editor, *Agreement Technologies*, volume 8, pages 233–249. Springer, 2013.
- [26] H. Perkins and A. Berkowitz. Perceiving the community norms of alcohol use among students: Some research implications for campus alcohol education programming. *International Journal of the Addictions*, 21:961–976, 1986.
- [27] B. T. R. Savarimuthu and S. Cranefield. Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent and Grid Systems*, 7(1):21–54, 2011.
- [28] P. Schultz, J. Nolan, R. Cialdini, N. Goldstein, and V. Griskevicius. The constructive, destructive, and reconstructive power of social norms. *Psychological Science*, 18(4):429–434, 2007.
- [29] P. W. Schultz, R. J. Bator, L. B. Large, C. M. Bruni, and J. J. Tabanico. Littering in context personal and environmental predictors of littering behavior. *Environment and Behavior*, 45(1):35–59, 2013.
- [30] S. H. Schwartz. A theory of cultural value orientations: Explication and applications. *International Studies in Sociology and Social Anthropology*, 104:33, 2006.
- [31] S. Sen and S. Airiau. Emergence of norms through social learning. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1507–1512, 2007.
- [32] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.
- [33] X. Wang, M. Maghami, and G. Sukthakar. Leveraging network properties for trust evaluation in multi-agent systems. In *Proceedings of the IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 288–295. IEEE Computer Society, 2011.
- [34] C. Yu, M. Zhang, F. Ren, and X. Luo. Emergence of social norms through collective learning in networked agent societies. In *Proceedings of the International Conference on Autonomous Agents and Multi-agent Systems*, pages 475–482, 2013.