Development and Validation of an Agent-Based Simulation Model of Juvenile Delinquency

Tibor Bosse, Charlotte Gerritsen, Michel C.A. Klein
Vrije Universiteit Amsterdam, Department of AI, Amsterdam, The Netherlands
{tbosse, cg, michel.klein}@few.vu.nl

Frank M. Weerman
Netherlands Institute for the Study of Crime and Law Enforcement
Amsterdam, The Netherlands
FWeerman@nscr.nl

Abstract

This paper describes the development and validation of a dynamic multi-agent model to simulate social learning of adolescence-limited criminal behaviour. The parameters of the agent model have been calibrated using real-world data that has been collected in a large study. In addition, a measure for correctness has been developed. The validation shows that the developed model predicts delinquency substantially better than a baseline model that only uses the delinquency of an agent in the previous year.

1. Introduction

The area of Criminology is a multidisciplinary field, which has as main objective to analyse criminal behaviour; e.g., [10]. As such, its main research goals are to predict in which circumstances which types of criminal behaviour occur. Since a substantial amount of crimes is performed by juveniles, an important challenge within Criminology is the analysis of the emergence of criminal behaviour during adolescence.

To address this challenge, several theories have been proposed within the criminological literature, that all depend on *social learning*: the idea that adolescents easily copy the behaviour of their peers. According to this view, the social network of an adolescent can be seen as a multi-agent system in which various interactions take place over time. One of the influential social learning theories is the Differential Association Theory by [18] which was later expanded by [5]. This (informal) theory states that behaviour is learned in interaction with others. We learn most from the people we are in close contact with, like parents and peers. A second important theory states that there are two distinct categories of antisocial behaviour and offending, namely *life-course persistent* and *adolescence*-limited [15]. Life-course-persistent anti-social behaviour is caused by neuropsychological problems during childhood that interact cumulatively with the criminogenic environments across development, which leads to a pathological personality. Adolescence-limited antisocial behaviour is caused by the gap between biological maturity and social maturity. It is learned from antisocial models that are easily mimicked, and it is sustained according to the reinforcement principles of learning theory. In this paper we focus on the second group, the adolescence-limited offenders, since their behaviour emerges from interaction with others.

When we take a closer look at the problem of the adolescence-limited offenders, several questions may be asked, among which:

- how does the delinquency level of adolescents relate to their personality traits?
- how do the delinquency levels of adolescents and their peers relate?
- how does the level of delinquency change over time?

To answer such questions, this paper proposes to make use of Agent Based Social Simulation (ABSS) techniques [7]. Since ABSS combines the advantages of the agent paradigm (e.g., personal characteristics of the individual agents) with those of social simulation (e.g., the possibility to perform scalable social "experiments" without much effort), it turns out to be particularly appropriate to analyse phenomena within the criminological domain [17]. Indeed, in recent years, a number of papers have successfully tackled criminological questions using ABSS, e.g., [4, 13, 14].

The current paper presents a multi-agent model that can be used to simulate the development of youth delinquency in a classroom, based on individual personality traits on the one hand, and the social

978-0-7695-3823-5/09 $26.00 © 2009 IEEE
DOI 10.1109/CSE.2009.136

network on the other hand. To calibrate the parameters of the model, data from an existing empirical study [21] have been used. In that study, the social networks of 1730 non-delinquent, minor delinquent and serious delinquent pupils at lower level secondary schools in the Netherlands were analysed. In addition, another dataset from that study (addressing a different set of schools than used for the calibration) has been used to validate the model. In future work, this model could be used to perform "what-if simulations" that can be helpful for policy makers, e.g., to investigate what is the best way to divide pupils over classes.

In Section 2 the empirical study on which our model is based is briefly summarised. Thereafter, in Section 3 we will introduce the overall approach that is used to develop and evaluate the model. Details of the simulation model are presented in Section 4, and in Section 5 validation of the model is discussed. In Section 6 related work is presented. Finally, Section 7 concludes the paper with a discussion and some ideas for future work.

2. Data Collection

The data used in this research come from the NSCR 'School Project' [21], a Netherlands based longitudinal study that focuses on peer network formation, personal development, and school interventions in the development of problem behaviour and delinquency. The sampling procedure was guided by two aims: one, to obtain a relatively 'high-risk' sample with a substantial proportion of delinquent young people, and two, to achieve enough variation in school contexts and student populations to be able to better generalise results. In order to achieve the first aim, schools and students in the lower educational strata of a major Dutch city with inner-city problems were over-represented. To achieve the second aim, students were also recruited from schools in smaller cities and towns in the vicinity. Although the sample is not a random sample, it can be considered representative of Dutch youths attending this school type (lower vocational) in the South West region of the Netherlands. In the whole country, 60% of young people attend this type of school.

For the current research, we used a cohort of students that started high school during the school year 2001/2002. The first year of secondary education in the Netherlands is comparable with 7th grade in the United States (most students are 12 or 13 years old). These students were surveyed during three consecutive years: 2002, 2003 and 2004.

Respondents' delinquent behaviour was measured using self-reports of a variety of offences. The self report method is a standard procedure in the field of Criminology, and it results in fairly reliable estimates of delinquency levels of young people, when it is conducted in a proper way and in an anonymous setting. Respondents were asked if they had ever committed an offence and, if so, how often during the reference period. The reference period covered the interval between the last summer holiday prior to the beginning of the school year and the time when the survey was administered (spring). The measures of self-reported delinquency used in this study come from 12 questions, among which: in the last year, how many times did you: "paint graffiti", "vandalise property", or "steal small things from shops worth less than 5 Euros" The total delinquency measure indicates how many types of these 12 delinquent behaviours were reported by the person.

The composition of *student networks* was studied using questions inspired by research carried out previously in this area. Respondents were provided with a numbered list of all students in their school year (so first-year students had the names from all fellow students in their own class as well as from all in the other first-year classes in their school). Then they were asked with whom they spent a lot of time (their school contacts – up to 10 fellow students could be identified, two of which could be labelled as "best friends"). In the analyses, friends' numbers were linked to the respondent's own number, enabling the networks of friends to be mapped and analyzed.

Apart from the central measures of delinquency and peer network composition, the study also used a substantial number of other measures on risk factors that are central in criminological theories and have been found to correlate with delinquency in the past. These risk factors are: low supervision by parents, low support by parents, low bond with school, low law conformity, high impulsivity, high adventure and risk-orientedness, high temper, much material needs, many time spent with friends, high deviance reinforcement by peers, being a member of troublesome youth groups. The relative position of students with regard to these risk factors were also obtained through the questionnaire: each risk factor is represented by a number of question items that were combined to scales (see [20] and [21] for more information).

3. Approach

As the goal of this research is a model that can be used to realistically simulate the development of juvenile delinquency in a multi-agent setting, we followed structured methodology to develop this model. As a first step, we built an initial dynamic model for the development of delinquency through social learning in a class room, based on an analysis of the literature. A comprehensive description of this step is provided in [2]. The model describes the influences of several personal characteristics, as well as the influences of other peers.

See Figure 1 for an overview: the box to the right depicts 'agent 1' (which represents a particular pupil). The delinquency of this agent is influenced by its previous delinquency (hence the circular arrow), its individual personality traits shown to the right (i.e., impulsivity, risk-orientedness, and so on, see previous section), and the external factors depicted to the left (i.e., the school, the parents, and a couple of peers, which are represented by a similar model as agent 1). Although the agents in the model are not very complicated, the multi-agent approach is an essential element for the simulations. The behaviour of each agent is influenced both by individual characteristics and by the relationship with up to 12 other agents in their network (which differs per agent).

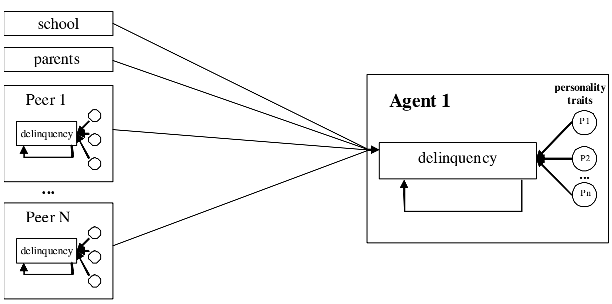

Figure 1. Overview of the simulation architecture.

The original model (from [2]) has the form of a set of differential equations. It has been shown that this model can be used to simulate delinquency development of a small set of agents in a classroom. The simulations exhibited several patterns that would be expected based on the criminological literature [2]. However, these initial simulations were not yet evaluated with empirical data, which is the main focus of the current paper.

Second, the dataset that is described in Section 2 has been split up in a training set and test set. Each set contains the data of around 250 pupils. When making this split, we guaranteed that there was no overlap between the schools used in the training set and those used in the test set. This way, we avoided that friendship relations exist *across* the boundary (i.e. that some pupils in the training set have friendship relations with pupils in the test set). Moreover, this approach makes it possible to detect whether the model has over-fitted to specific schools (which is not possible if the pupils used for the test set and the training set come from the same school).

Third, an evaluation measure has been developed that can be used to quantify the correctness of models and to discriminate between the accurate and less accurate models. This measure accommodates the intuitive ideas about a correct prediction in one number. The precise description of this measure is given in Section 4.1.

In a next step, we tried to calibrate the model with the data in the training set. This has been done by taking the model from [2] (extended with some additional factors reported in [21]) as a basis, and systematically adjusting it and comparing the agreement of the simulation results with the actual measurements in the training set. The adjustment consisted of both ignoring factors in the model (i.e. leaving out variables in the formulae) and calibrating parameters (i.e. changing the value of weighting variables). The idea behind this was to take a 'principle of parsimony' approach: although the original model was composed of factors that (according to the criminological literature) play a role in juvenile delinquency, this does not necessarily mean that the best model contains all of these factors. The aim of this phase was to achieve a list of models that resulted in a high correctness score.

Finally, the second data set was used to validate the different variations of the model that seemed promising during the calibration phase. In this phase, we did not change the model or parameters, but just calculated the accuracy according the developed measure for all formulae that resulted in a high score in the first phase. This method gives an unbiased validation of the accuracy, as the validation is performed on a different data set than the tuning.

4. Models for Simulating

Before development of a simulation model itself, a measure for evaluation has been established. This is described in the next section. The subsequent sections discuss the actual models and the calibration phase.

4.1. Evaluating Measure

The development of an evaluation measure is an important step, since it has a large impact on which models are considered accurate, and which are not. Intuitively, a simulation model for juvenile delinquency can be considered accurate when for a given set of pupils (and their personal characteristics) it predicts correctly when which pupil will show delinquent behaviour. To develop an evaluation measure, each pupil is assigned a delinquency score based on their answers to the questions related to their delinquent behaviour. However, for practical reasons, the 12-point scale used in the empirical study is converted to a binary scale: all pupils that have a delinquency score of ≥1 are assigned the value 1 (i.e., delinquent), and all other pupils are assigned the value 0 (not delinquent). The main motivation for this is that the distribution of the empirical data is not uniform: by far, most of the pupils have a delinquency score of <1 on the original scale, whereas only a few of them have a score of ≥1. For the simulated pupils (the agents) the same conversion will be made.

Next, the issue of time should be solved: for which moments will the simulated data be compared with the empirical data? For this, we simply used the time points for which the empirical data was available, i.e., after 1 year and after 2 years. Thus, the problem of defining a measure of evaluation has been reduced to matching the 0's and 1's of the pupils in the empirical data to the 0's and 1's of the corresponding agents in the simulation, both for the time point after 1 year and after 2 years. A good solution to this problem would provide points for *hits* (i.e., cases where both the data and the simulation result in 1) and *correct rejections* (i.e., data=0 and simulation=0), but may provide penalty points for *misses* (data=1 and simulation=0) and *false alarms* (data=0 and simulation=1). To select such a measure, it is worthwhile to consider the following standard measures from signal detection theory [11]:

Hit Rate = Hits / (Hits + Misses)
False Alarm (FA) Rate = FA / (FA + Correct Rejections)
Accuracy Rate = (Hits + Correct Rejections) / (Hits + Misses + FA + Correct Rejections)

For the current purpose, all four elements used in these measures are relevant: the model should maximise the amount of hits and correct rejections, and minimise the amount of misses and false alarms. Therefore, it makes sense to select the Accuracy Rate as evaluation measure. However, another problem is that in the empirical data, the amount of pupils that show delinquent behaviour is much smaller than the amount of pupils that do not show delinquent behaviour (ratio of about 1:3). Given this information, the extremely simple strategy to *always* predict non-delinquency would already result in a relatively high Accuracy Rate. To illustrate this, imagine a classroom of 100 pupils, of which 25 show delinquent behaviour and 75 do not. In that case, the above strategy would provide 0 hits, 75 correct rejections, 25 misses, and 0 false alarms, thus an Accuracy Rate of 75/100 = 0.75. This is not very desirable, since all crucial cases (the delinquent pupils) are missed here. To compensate for this, the following alternative Accuracy Rate has been defined:

Accuracy Rate 2 = (w*Hits + Correct Rejections) / (w*Hits + w*Misses + False Alarms + Correct Rejections)

Here, w is a weight factor, representing the importance of finding hits. For the current research, it makes sense to choose w=3, since in the data the ratio of delinquents vs. non-delinquents is more or less 1:3. Given this new formula and weight factor, the Accuracy Rate of the simple strategy defined above would become 75/150 = 0.5, which is more reasonable. For this reason, both during the calibration and the validation phase (explained below), this alternative Accuracy Rate has been used.

4.2. Development of the Simulation Model

In the simulation experiments, the delinquency of a large number of agents is calculated according to a set of formulae that determine how the delinquency of an individual agent is influenced by its own delinquency in the previous year, its personal characteristics, its peers and the delinquency of its peers. Eventually, this can be used to simulate the development of the delinquency of a large number of agents over several years.

To perform these simulations, we used standard numerical simulation software. The multi-agent system was modelled as a multi-dimensional array, where each array represented a different agent. The different dimensions represented characteristics of the agents over time. For example, these dimensions specified the individual characteristics (like impulsivity and risk-orientedness) and the relations to peers. To calculate the new delinquency of each agent, the following algorithm was used (in pseudo code):

```
For each agent:
1. determine current delinquency
2. determine individual characteristics
3. compose the social network (friends)
4. calculate average delinq. of social network
5. calculate new delinquency, using
   information from step 1, 2, and 4
```

For the calculation of the new delinquency of the individual agents (step 5), various variants of the model have been tried. The initial model that was used takes all risk factors mentioned in Section 2 into account for which was determined that they were statistically significant (all except "many time spent with friends" and "being a member of a troublesome youth group"). The data set contained scores for all these factors, but the maximum score for the distinct factors differed. Therefore, the scores first had to be scaled into the same dimension, i.e., between 0 and 1. As weighting factors, the delinquency *odds ratio's* for the individual risk factors were used. An odds ratio is defined as the ratio between the odds of an event occurring in one group and the odds of it occurring in another group. In this particular case, they provide information about the chance on delinquent behaviour when a certain risk factor (e.g., impulsiveness) is present [21]. The *main formula* in the initial model (i.e., the formula to calculate delinquency in the next year, step 5) consisted of a weighted sum of all risk factors (except the "deviance reinforcement by peers") and the average delinquency of friends in the previous year multiplied by the "deviance reinforcement by peers". The idea behind the specific treatment of the "deviance reinforcement" factor is that this factor will leverage the influence of peers.

The main formula of the initial model is the following:

$$
\begin{aligned}
&\text{delinquency}(y) = (\, w \cdot \text{delinquency}(y-1) + \\
&\quad \text{odds}_{\text{parent\_supervision}} \cdot \text{parent\_supervision}(y) \\
&\quad + \text{odds}_{\text{parent\_support}} \cdot \text{parent\_support}(y) \\
&\quad + \text{odds}_{\text{bond\_school}} \cdot \text{bond\_school}(y) \\
&\quad + \text{odds}_{\text{law\_confirmity}} \cdot \text{law\_confirmity}(y) \\
&\quad + \text{odds}_{\text{impulsivity}} \cdot \text{impulsivity}(y) \\
&\quad + \text{odds}_{\text{risk\_orientedness}} \cdot \text{risk\_orientedness}(y) \\
&\quad + \text{odds}_{\text{temper}} \cdot \text{temper}(y) \\
&\quad + \text{odds}_{\text{material\_needs}} \cdot \text{material\_needs}(y) \\
&\quad + \text{odds}_{\text{deviance\_reinforcement}} \cdot \text{deviance\_reinforment}(y) \\
&\quad * \text{average\_delinquency\_friends}(y-1)\,)\, / \, w + \textstyle\sum_{\text{odds}}
\end{aligned}
$$

This formula formed the basis for the calibrations in the next phase.

4.3. Calibration of the Model

To fine-tune the model, we followed a systematic approach in which we identified the contribution of all elements of the initial model to the overall correctness of the simulation results, by leaving out specific risk factors and / or changing the weighting or combination mechanism. In order to do so, we simulated the behaviour of delinquency for the pupils in the training set (which involved 194 pupils) and compared this to the actual data, using the evaluation measure as described in Section 4.1. This resulted in a correctness score for each variation of the model.

However, there is one additional concern. The simulation yields a *real* value for the delinquency, while the abstracted actual data is a binary value (0 or 1). To compare them, a specific threshold for the simulated delinquency has to be used above which an agent is considered to be delinquent. This threshold has a different optimal value for each variation of the model. Therefore, the correctness of a specific variant of the model is actually the highest accuracy for all values of the threshold. Thus, if the accuracy of a model variation is reported, the threshold at which this accuracy is reached has to be mentioned as well.

To be able to compare the quality of the predictions of the model with a baseline, we included four baseline predictions (variant 1-4). These predictions simply say, respectively, that all agents will become delinquent (i.e. regardless of all data, it will always output "1"), that no agent will become delinquent (always "0"), or predict a random distribution of delinquents and non-delinquents, either in the ratio 1:1, or in the same proportion as the empirical data(1:3). Next, different variants of the model were tried (variant 5-11), starting with the initial formula from Section 4.2. During the calibration, it turned out that many risk factors did not have a positive influence on the accuracy of the model. Leaving them out resulted in a higher accuracy rate than including them in the model. In addition, it appeared that the delinquency in the previous year is a very good predictor for the delinquency in the next year. In many variations of the model, we therefore used a disjunction of the delinquency in the last year and a combination of risk factors. Table 1 lists the most promising models, their accuracy according to the developed evaluation measure and the thresholds and weights at which this value was achieved.

Model variant 10 yielded the highest accuracy scores. This model uses, in addition to the delinquency in the previous year, the impulsivity and the product of the deviance reinforcement and the delinquency of the best friends. Conversations with criminological experts pointed out that these factors are well in line with the main theories in criminology. For the selected data of 194 pupils, model variant 10 resulted for the simulation of the first year in 46 hits, 44 false alarms, 11 misses, and 93 correct rejections, and for the second year in 46 hits, 43 false alarms, 8 misses, and 97 correct rejections. All other variations resulted in lower accuracy scores.

In addition to the above described evaluation measure, we also assessed the quality of the model using an ROC curve analysis [9]. An ROC (Relative

Table 1. Variants of the model and accuracy values.

Model variant	Main formula	Optimal threshold	Weights	Accuracy y1 → y2	Accuracy y2 → y3	Average accuracy
1	always 0			44.48	46.36	45.42
2	always 1			55.52	53.64	54.58
3	random 0 and 1 (with ratio 1:1)			51.95	48.68	50.31
4	random 0 and 1 (with ratio 3:1)			51.95	44.37	48.16
5	initial formula (see Section 4.2)	0.59		62.99	69.21	66.10
6	delinquency(y-1)			67.21	73.84	70.52
7	delinquency (y-1) OR delinquency_friends(y-1)	0.14		65.58	76.49	71.04
8	delinquency (y-1) OR delinquency_friends (y-1) OR risk_orientedness(y) > w1 OR temper(y) > w2	0.34	14, 12	70.45	74.83	72.64
9	delinquency (y-1) OR (odd * risk_orientedness(y) + odd * deviance_reinforment(y) * delinquency_friends(y-1)) / $\sum_{odds}$	0.55		72.08	75.17	73.62
10	delinquency (y-1) OR (odd * impulsivity(y) + odd * deviance_reinforment(y) * delinquency_best_friends(y-1)) / $\sum_{odds}$	0.31		75.00	77.81	76.41
11	delinquency (y-1) OR (odd * impulsivity(y) + odd * deviance_reinforment (y) * (w * delinquency_ friends (y-1)) + (1-w) delinquency_best_friends (y-1)) / $\sum_{odds}$	0.31	0.35	75.32	77.15	76.24

Operating Characteristic) curve is a graphical plot of the fraction of true versus the fraction of false positives for a binary classifier system as its discrimination threshold is varied. The threshold in our model is the value of the calculated delinquency above which a pupil is classified as delinquent.

Figure 2 shows the graph of the ROC curve for the best model variant, number 10. We also calculated the area under the ROC curve (AUC), a scalar measure for the quality of the predictions. For model variant 10, the AUC is 0.79. An AUC-value larger then 0.70 is called ‘acceptable’, larger then 0.80 ‘excellent’ and larger then 0.90 ‘outstanding’ [12].

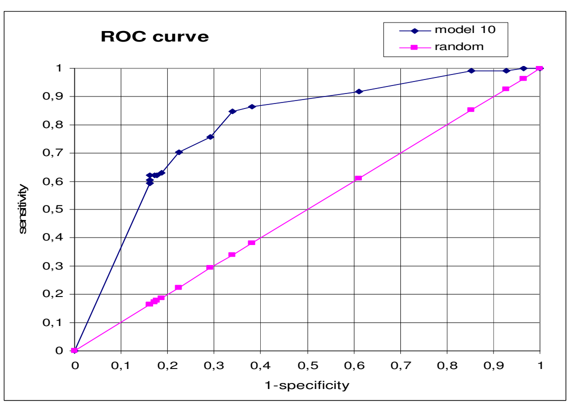

Figure 2. ROC curve for model 10 and a random prediction.

5. Validation

In order to validate the models presented in the previous section, a second dataset was used. Like the dataset used for calibration of the model, this second dataset was also taken from [21]. Thus, for each pupil the same types of information (i.e., delinquency measures, peer networks, and individual risk factors) were available; only a different pool of pupils was taken. However, as mentioned earlier, we guaranteed that there was no overlap between the schools used in both datasets. This second dataset involved 299 pupils.

As mentioned earlier, for the validation the same formulae as used for the calibration were used. This means that the same parameter values, (e.g. for thresholds, weight factors) were used. To be able to evaluate the results of the different models, also the baseline models were applied to the validation dataset. The empirical data for this dataset (of 299 pupils) were fed as input to the different models, and the Accuracy Rate was calculated.

The results of applying the different models to the validation dataset are shown in Table 2. As shown by this table, the more sophisticated models (i.e., variant 8-11) are clearly more accurate (varying from 65.35 to 66.33) than the baseline strategies (varying from 41.84 to 58.16), the initial model (57.53) and the model that predicts stability with respect to the previous year (60.53). Surprisingly, for this dataset the strategy of taking the peer network into account (variant 7) does not seem to add much with respect to the strategy of looking at the previous year only (variant 6), but does make a difference when taken in combination with the deviance reinforcement factor (later variants).

Furthermore, it is worthwhile to note that most overall accuracy rates are slightly lower than they were for the first dataset. This is the case for all variants, except for model 1-4 (which is obvious, since these models do not make use of any predicting factor). It is however surprising that also variant 6 (the measurement that only takes the previous year into account) scores lower here than for the first dataset. This in an indication that this second dataset was simply less ‘stable’ than the first dataset: there were more changes in delinquency, which makes it more difficult for a model to make accurate predictions. Despite this more difficult dataset, the performance of the best model (variant 8) was still more than 6 points better than the performance of the straightforward strategy of variant 6, which is about the same difference as was found for the first dataset.

Table 2. Validation results.

Model	Acc. y1 → y2	Acc. y2 → y3	Avg. acc.
1	42.33	41.36	41.84
2	57.67	58.64	58.16
3	50.44	52.18	51.31
4	44.97	49.74	47.36
5	55.73	59.34	57.53
6	61.02	60.03	60.53
7	60.85	60.03	60.44
8	68.78	63.87	66.33
9	65.26	67.19	66.22
10	66.14	66.14	66.14
11	64.73	65.97	65.35

For the validation, we also plotted the ROC curve and calculated the AUC value. Similar to the results using the other evaluation measure, the AUC for the validation data was less then the AUC for the training data, i.e. 0.68. This is still much larger than the 0.50 that a random prediction would yield.

6. Related Work

With respect to related work, there are both commonalities with the social and behavioural sciences, and the AI and Computer Science.

Concerning the first, the current paper is related to important articles from the 1960’s and 1970’s such as [5, 18], which were the first to formulate (different variants of) the social learning theory. In fact, these theories formed the basis of the research questions addressed in this paper. Based on these theories, [16] identified a number of (informal) properties that are expected to hold for social learning in Criminology. The simulation model presented within this paper indeed satisfies these properties. Next, a number of papers in Criminology propose more refined models for social learning, often focusing on specific aspects of the learning. For example, [19] compared three theoretical (but not computational) models of the interrelations among associations between delinquent peers, delinquent beliefs, and delinquent behaviour. Finally, several authors have performed empirical studies on social learning of delinquent behaviour in schools, e.g., [20]. Our model was designed explicitly with the purpose of reproducing such data.

Concerning the literature in AI and Computer Science, we are not aware of approaches using multi-agent technology to simulate delinquent behaviour of individuals in a group. However, various papers have similarities to the work proposed here. First, [8] presents a model that uses differential equations to describe the development of juvenile criminal behaviour. They aim for an integration of multiple criminological theories, whereas we focus (in more detail) on the former only. Moreover, several authors have created models that address social learning and criminal behaviour at a more global level e.g. [22]. These models differ from our model in the sense that they are situated at a macroscopic level, thereby abstracting from differences between individuals. Furthermore, a large number of approaches address simulation of the environmental aspects of criminal behaviour, such as the displacement of crime and the emergence of “hot spots”, e.g., [1, 13]. Finally, relevant work is put forward by [6]. They identify a number of (cognitive) factors that are relevant in social learning in general. However, in contrast to our work, they do not provide a computational model.

7. Conclusion

This paper contributed the development and validation of a dynamic agent-based approach to simulate social learning of adolescence-limited criminal behaviour. This approach has been used to perform simulation experiments in which the delinquency of 250 pupils is dynamically calculated over a couple of years. This expected delinquency is based on personal characteristics on the one hand and the delinquency of peers on the other hand. A second dataset has been used to validate the model, using a specifically developed accuracy measure. The validation shows that the model predicts delinquency substantially better than a baseline model that only uses the delinquency of the previous year.

Note that an inherent consequence of the use of empirical data is that such data is often incomplete. This incompleteness may be caused by respondents not answering all the questions or by the fact that some respondents reported friends that were not part of the study. Such incompleteness is one of the complicating factors in the development of an accurate simulation model. However, an important advantage of the approach presented in this paper is that it does not use one single formula to calculate future delinquency, but presents a whole range of different formulas. As a result, for a particular real-world case, the modeller can choose the particular formula that best fits the available information. E.g., if no information about the participant's impulsiveness is available, then a formula can be selected that does not make use of this factor[1]. This property of flexibility (and user transparency) of the model is an important advantage over, e.g., approaches based on machine learning. Nevertheless, for future work it is worthwhile to explore whether automated learning techniques can be exploited to improve (at least parts of) the model.

As soon as the model is sufficiently validated, an interesting direction for future work is to perform so-called "what-if simulations", or computer-supported thought experiments. These thought experiments can be particularly useful in policy making. An interesting question could be, for example, "what would happen if we placed one bad child in a classroom full of teacher's pets"? Will this delinquent pupil adapt himself to the environment and become good as well, or will he manage to make the entire class a bit more delinquent? How will the average class level evolve? The answers to these questions may be very important for high schools, to decide how to fill in their classes. In future work, it is planned to perform a number of such what-if simulation in a systematic manner, in collaboration with experts from Criminology.

References

[1] Bosse, T., and Gerritsen, C. (2008*). Agent-Based Simulation of the Spatial Dynamics of Crime: On the Interplay between Criminal Hot Spots and Reputation.* In: Proc. of the 7th International Joint Conference on Autonomous Agents and Multi-Agents Systems, AAMAS'08. ACM Press, 2008, pp. 1129-1136.

[2] Bosse, T., Gerritsen, C. and Klein, M.C.A. (2009). *Agent-Based Simulation of Social Learning in Criminology*. In: Proc. of the Int. Conf. on Agents and AI, ICAART'09. INSTICC Press, 2009, pp. 5-13.

[3] Bosse, T., Jonker, C.M., Meij, L. van der, and Treur, J. (2007). A Language and Environment for Analysis of Dynamics by SimulaTiOn. *International Journal of AI Tools*, vol. 16, issue 3, pp. 435-464.

[4] Brantingham, P. L., and Brantingham, P. J. (2004). Computer Simulation as a Tool for Environmental Criminologists. *Security Journal*, 17(1), 21-30.

[5] Burgess, R., and Akers, R.L. (1966). A Differential Association-Reinforcement Theory of Criminal Behavior. *Social Problems*, vol. 14, pp. 363-383.

[6] Conte, R., and Paolucci, M. (2001). Intelligent Social Learning. *Journal of Artificial Societies and Social Simulation*, vol. 4, issue 1.

[7] Davidsson, P. (2002). Agent Based Social Simulation: A Computer Science View. *Journal of Artificial Societies and Social Simulation*, 5(1).

[8] Dijkum, C. van, and Landsheer, H. (2000). Experimenting with a Nonlinear Dynamic Model of Juvenile Criminal Behavior. *Simulation & Gaming*, vol. 31, pp. 479-490.

[9] Fawcett, T. (2006). An introduction to ROC analysis. Pattern Recognition Letters, 27, 861-874.

[10] Gottfredson, M. and Hirschi, T. (1990). A General Theory of Crime. Stanford University Press.

[11] Green, D.M., Swets J.A. (1966) *Signal Detection Theory and Psychophysics*. NY: Wiley.

[12] Hosmer, D., S. Lemeshow (2000). Applied logistic Regression. New York, John Wiley & Sons Inc.

[13] Liu, L., Wang, X., Eck, J., and Liang, J. (2005). Simulating Crime Events and Crime Patterns in RA/CA Model. In F. Wang (ed.), *Geographic Information Systems and Crime Analysis*. Singapore: Idea Group, pp. 197-213.

[14] Melo, A., Belchior, M., and Furtado, V. (2005). Analyzing Police Patrol Routes by Simulating the Physical Reorganisation of Agents. In: Sichman, J.S., and Antunes, L. (eds.), *Multi-Agent-Based Simulation VI, Proc. of the 6th Int. Workshop on Multi-Agent-Based Simulation, MABS'05*. LNAI, vol. 3891, Springer Verlag, 2006, pp 99-114.

[15] Moffitt, T.E. (1993). Adolescence-Limited and Life-Course-Persistent Antisocial Behavior: A Developmental Taxonomy. *Psychological Review*, 100(4), pp. 674-701.

[16] Opp, K.D. (1989). The Economics of Crime and the Sociology of Deviant Behaviour - A Theoretic Confrontation of Basic Propositions. *Kyklos*, vol. 42, issue 3, pp. 405-430.

[17] Sun, R. (ed.) (2006). Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation. Cambridge University Press.

[18] Sutherland, E.H., and Cressey, D.R. (1966). *Principles of Criminology*, 7th edition. Philadelphia: J.B. Lippincott.

[19] Thornberry, T.P., Lizotte, A.J., Krohn, M.D., Farnworth, M., and Jang, S.J. (1994). Delinquent Peers, Beliefs, and Delinquent Behavior: A Longitudinal Test of Interactional Theory. *Criminology*, vol. 32, pp. 47-83.

[20] Weerman, F.M., and Bijleveld, C.C.J.H. (2007). Birds of Different Feathers. *European Journal of Criminology*, vol. 4, issue 4, pp. 357-383.

[21] Weerman, F.M., Smeenk, W., and Harland, P. (eds.), i.c.w. Ezinga, M., Slotboom, A.-M., Bijleveld, C., Laan, P. van der, and Westenberg, M. (2007). *Problem behavior of students during secondary education: Individual development, student networks and reactions from school* (in Dutch). Amsterdam: Aksant.

[22] Winoto, P. (2002). An Agent-Based Simulation of the Market for Offenses. In: *AAAI Workshop on Multi-Agent Modeling and Simulation of Economic Systems*. Edmonton, Canada.

[1] Of course, there should be at least some factors for which information is available, but this holds for any predictive model.