

Simulating the dynamics of organized crime using a multi agent system

Manon de Vries
Rijksuniversiteit Groningen, The Netherlands
manon.de.vries.41@student.rug.nl

Bart Verheij
Rijksuniversiteit Groningen, The Netherlands
b.verheij@ai.rug.nl

ABSTRACT

In criminology, two types of organized crime are found: a hierarchical type, with clear roles and a pyramidal structure (traditional mafia), and a more flexible type, with less hierarchy and flexible configurations. By making a multi agent model I tried to gain insights in how the cooperations work in these criminal organisations, and particularly on what causes the two different types. Herefor three factors are investigated: the role of the initialisation of the model, the role of the amount of experts and the influence of the cooperation threshold (how critical agents are when choosing their colleagues). The results do not show a clear relation between the initialisation and the amount of experts on the hierarchy of the network. However, the segmentation of the network does seem to depend on the cooperation threshold. This could point at the fact that a different level of distrust may be the cause of the two types of organized crime.

INTRODUCTION

Experiments in criminology are difficult [7]. They are often extremely time-consuming (e.g. looking at criminal trajectories), not ethically possible (sending one thief to prison and another home) or costly. With the scarce amount of experiments, findings from observations and information from lawsuits, theories are formed. Unfortunately, these theories can not be tested to their fullest extent. Most of the time a consensus is reached about the validity of a theory, but the underlying mechanisms and the applicability are often unclear.

In this article we want to introduce a partial answer to this problem. We can shed some light on underlying mechanisms and dependability on different parameters. This can be achieved by combining a different research domain with criminology: artificial intelligence. By using multi-agent models, we can try to reproduce (criminal) behaviour computationally. We implement underlying causes suggested by theories into agents and design a relevant environment for them to interact in. When a proper model is made, with correct behaviour rules for the agents, the same emergent behaviour as in real life should occur.

The process of making such a model may, depending on its resemblance with reality, result in three useful insights:

1. When the predicted underlying causes do not produce the expected behaviour, the interactions are too simple or the underlying process suggested by literature is

different in real life. Although this is generally not the aim, it can still tell us a lot about our theory, by stressing its application [4].

2. When the model does replicate the results using the suggested underlying mechanisms, this can serve as extra (indirect) evidence for the theory.

When the model seems to be valid, a lot of new applications emerge. First the model can be used for virtual experiments. In the real world it is often not possible to test different policies on different people, but there are no ethical limitations when using virtual agents. Of course, we are talking about qualitative insights here, not quantitative, and we must be careful not to rely too much on these models. We can not conclude from only a simplified model that policy 1 will be 100 million euros cheaper and reduce criminality by 10 percent, but we can study the impact of policies on the dynamics of a system. This brings us to our next benefit:

3. Insight on how the criminal world works. The models can be used for explaining complicated theories by using simplification, interactivity and visualisation. One day these models may even help policy makers, who do not always have a full background in criminology, to make better-founded decisions.

By making models like this, a lot of knowledge and understanding is gained about the subject. Not only by the result, but particularly by the designing process. Programming forces the theorist to be explicit and to think about the details.

In this article we will stress the importance of the cooperation between criminology and artificial intelligence. Although computational models are used in the social sciences (for a nice comparison with computational sciences see Marks [8]), the cooperation with criminology is relatively rare (for a nice exception see [6]).

We will try to illustrate the benefits of this cooperation with a multi agent model of organized crime. For this we will first briefly discuss the developments in criminology, specifically in organized crime research and then we will introduce the AI-side; multi-agent systems and its applications.

RESEARCH CONTEXT

Organized crime

The view on criminology, specifically in the area of organized crime, has been undergoing a paradigm shift in the last decades. In the past, organized crime was found to be hierarchical, with a relatively stable configuration of members. It was seen as a domestic problem, troubling a relatively small number of states such as Italy, the United States, and Japan [16].

Since then the field of criminal research evolved and eventually a shift occurred, towards a more project-based view on criminality. The belief that there is a clear line between the upper and underworld was abandoned and a network-perspective gained importance (see for a relevant article van Calster [13]).

In these years not so much the organized crime changed, as the point of view of the researchers, shown for example in the article from Block [1]. He was one of the first to describe crime with a network-perspective, analysing the cocaine branch in early twentieth century New York. He analysed that even back then there was not one big organization, and that the existing small groups were highly flexible. Members had multiple roles, switched groups or were dealers in different configurations.

The big mafia-like organizations most of us think of when talking about organized crime are therefore only part of the picture.

Still, the old-fashioned ‘mafia’ view, based on countries like America and Italy [11], was used for many years after Block’s work.

In the ‘80s organized crime became a priority to politics in the Netherlands. Police began to judge criminal organizations on criteria which were based on this Italian-American view. Although not many organizations were found who met this criteria, there was a lot of media attention [9]. A number of incidents caused a feeling that the magnitude and problem of organized crime in Holland could be much larger than first anticipated. Multiple measures were taken based on this idea.

In the ‘90s the government composed a research group to get to the bottom of how organized crime was present in the Dutch society. The results nuanced the image of an octopus-like organisation. From this commission the monitor organized crime Holland was formed to better study organized crime in Holland in the following years. This monitor has so far reported about 120 criminal cases dating from 1996 till 2007, and is still monitoring cases today. Their information will be vital for this paper, as details of their work will be used later. One of their conclusions made in 2007 is the following:

In the Netherlands, pyramidal structures with a strict hierarchy, a clear division of tasks, and an internal sanctioning system, are the exception rather than the rule. In many cases of the Organized Crime Monitor the term ‘criminal networks’ is far better suited for describing the actual structure of cooperation. Offenders cooperate in certain projects, yet the structure of cooperation is fluid and changes over time.[2]




Figure 1: Simulation of Human Behaviour during Emergency Evacuations [10]

This was a relief for the politicians, because they feared organisations of Italian proportions in Holland.

Despite of this relief, the kind of criminal networks in the Netherlands also poses difficulties [2]. The dynamicality of these small, more segmented organizations makes fighting them a lot harder. It is important to find vulnerabilities and weak links to fight organized crime effectively [14][16]. Understanding the structure of these kind of organizations is therefore essential, the monitor says:

By analysing criminal associations as criminal networks, not only is it possible to get a clear view of the stability in certain associations, but also of discontinuity and change. Such insights are vital to police investigations.[2]

Multi-agent systems

The AI-research field we want to use to aid our criminological understanding of these networks is multi-agent systems. Multi-agent systems are computer simulations, consisting of multiple components, called the agents (which can be people, birds, ants, etc.), in this case the criminals. These agents have simple rules but can, by interacting together, create complex behaviour [17].

The agent-based computational model is a new scientific instrument. ‘...’ The central idea is this: To the generativist, explaining the emergence of macroscopic societal regularities, such as norms or price equilibria, requires that one answers the following question: “How could the decentralized local interactions of heterogeneous autonomous agents generate the given regularity?”[4]

Multi-agent models are used in many different areas, ranging from predicting crowd behaviour during emergency situations to modeling swarm behaviour of birds or fish (see Figure 1). The procedure is to situate a population of autonomous heterogeneous agents in a relevant (spatial) environment. Then allow them to interact according to simple local rules, and thereby generate - or “grow” - the macroscopic regularity from the bottom up.

Agent simulations possess some distinctive qualities:

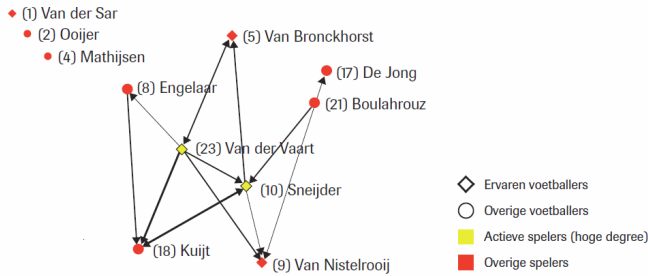


Figure 2: Degree-centrality in the dutch football team [15]. An example of a network view

- Agent populations are heterogeneous; individuals may differ in multiple ways - genetically, culturally, or by preferences - all of which may change or adapt over time.
- They are autonomous; there is no central, or “top-down”, control over individual behaviour in agent-based models. Often there will be feedback from macro structures to micro structures, but as a matter of model specification, no central controllers or other higher authorities are implemented.
- Important is that the notion of “local” is well posed. Events typically transpire on an explicit space, which may represent geographical environments. The environment may also represent another feature, like ‘knowledge space’ [7]. In these spatial models, the agents have coordinates to indicate their locations. There is however a third option, which we will apply: to have no spatial representation but to link agents together in a network (see Figure 2). In this case the location only has a relational meaning, like centrality in the network (see for more about this Scott [12]). The coordinates of the agents are calculated by using one of many placement algorithms which shows the agents in a way that is optimal and easy to distinguish, for example, by displaying disjoint groups separately.
- Agents behave according to simple local interactions. Typically, agents interact with neighbours in the environment (and perhaps with environmental sites in their vicinity). Uniform mixing is generically not the rule [4].
- In most of multi agent systems, agents have limited information and limited computing power.

Model assumptions

Let us look back. Organized crime in Holland is project-based and flexible. Because of the flexibility, it is harder to fight it, as pointed out by the monitor organized crime [2]. Understanding these networks is therefore of key-importance. Multi-agent systems provide a tool to analyse social networks. These models can tell us something about how the members of a network cooperate.

We want to implement the characteristics known of the members of criminal networks, and see if these characteristics can

explain the more segmented, flexible, projectbased nature of organized crime in the Netherlands and the hierarchical organized crime in Italy and America. The results may tell us about the factors involved. The overarching goal is to provide some insight in the underlying structures and dynamics of the complex relations in organized crime.

The configurations are fictive and not based on an existing network, but the rules are deduced from various articles which describe organized crime [2][3][9][1]. Our most important source is the Dutch report on organized crime [2], made by the monitor organized crime Holland, as mentioned earlier. Their longlasting investigation describes criminal careers, social embedding of organized crime, relations between members, structure of networks and internationality of networks. It describes both the macro structure - flexible configurations, project crimes - as the characteristics of individuals, like the importance of trust, the different role of facilitators and the parameters which seem to influence the chance of cooperation (capacity of a person, trust and expertise).

Three important findings of the committee were incorporated in the model:

1. “The fulfilment of an illegal transaction is naturally not forcable by a law. That is why offenders - ironically - have to rely strongly on reliable partners.” [2, Author’s translation]
 2. “In many cases, there are clear key players on whom many other offenders depend because of their financial resources, know-how or contacts. ‘...’ Little by little, other offenders may become less dependent of these key players, as they gather money, know-how and contacts themselves and subsequently start engaging in their own criminal activities.” [2]
- Note here the importance of reliability and the importance of the criminal’s capacities (how much can you produce, deliver or sell), and that these can change over time.
3. “It may also become clear that facilitators, who usually operate in the periphery of criminal organizations (such as money changers, underground bankers, financial service providers and forgers), render services to several criminal organizations. This way, they occupy a more important position within a criminal network than the consideration of separate ‘criminal organizations’ would make one think.” [2]

We incorporate this idea in the model by simulating two types of agents; experts (facilitators and normal agents).

Hypotheses

We will investigate three hypotheses:

1. With a hierarchical initialisation, the resulting network will be more hierarchical than with a random or flat initialisation.
In the model we will always repeat the experiments with three different initialisations. This initialisation affects who knows who before the simulation starts.

Members can only cooperate with people they know, so the initialisation should influence who connects with who. This will be measured by the average hierarchical value of the agents (which will be explained later).

2. When the agents are being more reserved/more picky in accepting new colleagues, as modeled by a higher initial cooperation threshold, this will result in a more segmented network: More groups, with less members and more dying of agents.

The threshold of an agent will determine if the agent, when requested, accepts to cooperate. This threshold is compared to the value of the requester (which depends on its reliability and capacity). When the initial threshold is higher, less agent will become colleagues, causing some agents not to cooperate at all. These will disappear from the simulation. This effect will be greater with a higher threshold, resulting in a lower average of agents in the simulation (multiple can disappear in one time-step, but only one agent is added). Lastly, because the agent will cooperate less, the chance of one close-knit group will be smaller, and less networking will occur because of less cooperations causing the chance to drop that two groups join, so we expect more (but smaller) groups.

We hope to see less hierarchy as well, but because of the effect that probably only agents with high reliability and capacity will be successful (work together) with high initial threshold, we do not expect a big difference in hierarchy.

3. Experts ('facilitators') have a higher initial threshold, so when there are more experts, the network will be more segmented as well.

According to literature, facilitators are persons with expertise, of which there are not many [2]. Every (criminal) group needs a facilitator, which gives the expert an advantage. Because people are more willing to cooperate with them, they can choose more carefully who to work with. This is implemented by a higher initial cooperation threshold. As we expected in the second hypotheses, a higher threshold will result in a more segmented network. Normally experts are only a small percentage of the network, so their influence on the overall network will not be enormous. But if we increase the amount of experts, we expect this influence to grow with them, and cause the overall network to be more segmented.

METHODS

The simulation revolves around the connections between agents; who knows who, but most importantly, who works with who. Agents choose to cooperate with each other based on the trust and the capacities of the other agent. This rule is the core of the simulation. We will now explain the model in more detail.

Every agent has the following characteristics:

- Type. There are two types of agents, experts and normal agents.
- Position. Every agent has a location in the available space. This position is relative, and not meaningful in

any other way than to display the agents in relation to others. A placement algorithm [5] is used to display the agents in an optimal way, letting connections cross as little as possible and showing different groups separately.

- Reliability. Reliability is a number between zero and one which represent the reliability of the person; zero is unreliable, one is as reliable as possibly can. We chose for an upper border because in real life this seems to be the case also; for two people who never do unreliable things, you cannot say that one is more reliable than the other.
- Capacity. Capacity is represented by a number between zero and infinity. It stands for the amount an individual can offer, not specifying in which quantity, this can be products, services or other means for which agents need each other.
- Acquaintances list. This list contains agents the agent knows, meaning he can access the reliability and capacity of that agent in order to maybe cooperate with one of them.
- Colleagues list. This list contains all agents the agent currently cooperates with.
- Lonelytime. This parameter is zero when the agents cooperates with one or more agents. When he does not have any colleagues it counts the epochs (time steps) since the last cooperation.
- Cooperation threshold. Each agents has a threshold above which he will accept to cooperate with an agent when requested.

Each epoch each agents looks at his list of acquaintances. He asks the agent with the highest value (reliability*capacity) to cooperate with him. Next he looks if he has any cooperation request himself. If any, he picks the one with the highest $r*c$ and checks if this number is above his cooperation threshold. If so, he accepts the request and adds the requesting agent to his list of colleagues. In the last step the agent checks his current cooperations. If one of his colleagues value has dropped below his threshold, he ceases the cooperation and removes this agent from his colleagues list. When an agent is added as a colleague, the agent also inherits this new colleagues coworkers as new acquaintances. This is 'networking'; an agent can meet other agents through mutual colleagues.

The parameters of each agent are updated each epoch. Capacity is raised by 0.001, reliability by 0.0005 and the cooperation threshold by 0.0015 (these numbers were determined by fine tuning the model). Each epoch, some agents are randomly picked (chance of this happening to an agent is 0.02) and his reliability is lowered by half, as to simulate snitching, betrayal, or other activities which damage trust.

Each epoch the number of agents in the simulation can change. During the simulation, the lonely time of an agent increases when this agent does not cooperate with anyone. After it reaches 10, the agent will disappear from the simulation, simulating the disappearance of an offender from the




Figure 3: The three initialisations; random, even and hierarchical respectively.

illegal world.

If the population of agents is less than the maximum number of agents (default 100), a new agent is added to the world. This new agent will get some connections (acquaintances) with a few agents in the simulation. Which agent he knows is always determined randomly. Important to note is that only one agent is added per epoch, while several can disappear in one.

The ‘type’ is a boolean which determines if an agent has expertise. It is initialised randomly according to a predefined percentage of experts (default is 10%). Reliability is randomly initialised between 0 and 0.8 and capacity with a poisson distribution (with $\lambda = 4$). The position of an agent is determined by a force-directed algorithm [5], which uses the connections between agents for optimal placement. Lonelytime is initialised at zero.

The initial acquaintances of agents can be initialised in three different manners.

The first initialisation possibility is random (Figure 3a). Here two agents are picked, and there is a chance (in the model 2%) that a connection is created between them (connection meaning that they are eachothers acquaintances/know eachother). The result is an average of 2 acquaintances per agent (but some will have more and some will have none). The second initialisation is the even/flat initialisation. The goal here is to give each agent exactly the same number of acquaintances (in this example 4, see Figure 3b, in the simulation 2 per agent).

The hierarchical view uses an extra factor to determine the relations: the value of the agent, or its reliability * capacity. The ‘godfather’ idea is that the big boss is the best (highest value), and is connected with good colleagues (fairly high value), who are connected with people with again somewhat lower value, etc. We implemented this in a kind of tree structure (see Figure 3c). The leafs are the people with low value. Every agent (except for the leafs) is connected with three others (two with lower value, one with higher value). Because of the existence of the leafs (which have only one connection), the average amount of connections is around two. Interesting, increasing the branching factor will not raise this average, because it also creates more leaves (the mathematical proof of this is outside the scope of this article).

Visualisation

The visible result is seen in Figure 4. It consist of three parts. The left part allows you to choose the mode (play, several experiments), initialisation type (random, hierarchical or even) and adjust the parameters (disabled in experi-

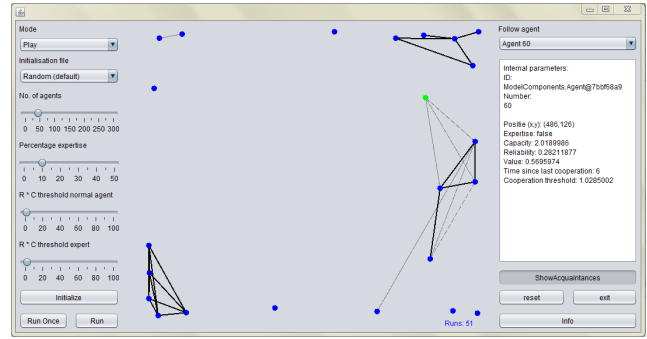


Figure 4: The visualisation with acquaintances

ment mode). The middle part shows a visualisation of the criminal network(s). The bold lines are cooperations, and when ‘show acquaintances’ is pressed, a grey, dotted line will appear between acquaintances. At the right side there is also the possibility (in play mode) to monitor an agent and follow it throughout the visualisation. This agent will appear green in the visualisation. While not selected, normal agents are shown blue and experts red.

Experiments

First a quick parameter search has been conducted to determine default settings for the network. The parameters are adjusted in a way that the simulation produces meaningful results; for example by making sure the thresholds are not too high, which would cause no agent to work with anyone, but also not too low, causing everyone to work with everyone. Except if mentioned otherwise, the following settings are used: percentage agents which are expert = 10 %, threshold normal agents = 1.25, threshold experts = 2, initial number of agents (and maximum during simulation) = 100, number of epochs per run = 1000. The average number of agents during the simulation will be lower, as mentioned before, because multiple agents can disappear in one epoch, while only one new agent is introduced per epoch. Each experiment is repeated 3 times for the three different initialisation, with the same random seed. Lastly, every experiment is repeated 20 times.

During the simulation, the following data is stored for analyses: the used initialisation (even, random, or hierarchical), the reliability*capacity value of the agents, the hierarchical value of the agents, the amount of agents in the simulation, the number of groups (where a group is defined as 4 or more agents working together), and the number of members of each group. When a parameter is changed during the experiment, this number is stored as well.

The hierarchical value needs some explaining. We needed something to measure the ‘hierarchy’ of the network. This is why we gave every agent a hierarchical value. This is defined as his own R*C value, plus the mean of the R*C values of his colleagues. This way, agents with a high reliability and capacity and colleagues with high reliability and capacity will score high. This value is then averaged over the agents in the simulation. This way, when agents cooperate with people of their niveau (similar R*C value), which we defined as hierarchical, the mean hierarchical value will be higher, than when they cooperate with all kind of people, of different niveaus, as happens more when there are

small groups (the agents have less people to choose from for a cooperation request, resulting in cooperations with more difference in value).

The role of initialisation in hierarchicality

This is the simplest of the experiments. It uses all the parameter setting mentioned before, only the initialisation file is changed. The resulting data file is analysed by a matlab script, which averages, summarizes and visualizes the data.

The influence of the initial cooperation threshold on the segmentation of the agents

To say something meaningful about the role of the influence of the cooperation threshold, we decided not to just increase the threshold of both experts and normal agents simultaneously with the same amount, or to change them separately, but to vary them both and report the output of every combination. We tested both with the values 0.25 - 2.5, with steps of 0.25. Above 2.5, no connections are formed anymore, and below 0.25 everyone connects with everyone. This results in $10 \times 10 = 100$ experiments with different parameter settings, done three times for the three initialisation files, and everything is repeated 20 times as mentioned before. The fact that we only repeat 20 times is caused by this experiment. The resulting datafiles are of such size, and the matlabscripts to analyse the data take so much time for the importing of the data, that more experiments would ask for resources which are not available to us at this moment.

The influence of the amount of experts on the segmentation

Using the settings mentioned in the beginning, we only vary the amount of experts. We test the percentages: 0 - 50, with steps of 5, resulting in 11 different experiments, conducted 3 times with the different initialisation files. This is again repeated 20 times.

RESULTS

The role of initialisation on the hierarchical value of the network

Our hypotheses was that when using the hierarchical initialisation, the resulting network would be more hierarchical than with the other two initialisations. As seen in Figure 5, a difference is seen between the three initialisations. The hierarchical initialisation does lead to more hierarchical networks. This difference is not big, and not significant when comparing it to the even initialisation (two-sample $t(38)=2.0, p=0.053$), but significant when comparing it to the random initialisation (two-sample $t(38)=3.0, p=0.004$). We can therefore conclude that the hypotheses is only partly true: the hierarchical initialisation does produce a more hierarchical network than the the random initialisation, but not more as the even initialisation.

What is interesting, although not within the scope of this experiment to explain, is that the hierarchical initialisation also produces groups with much more members (as seen in Figure 6), with an average of 7.3 compared to 4.1 for both random and even initialisation (compared to random: two-sample $t(38)=4.7, p < 0.001$, and compared to even initialisation: two-sample $t(38)=4.8, p < 0.001$). Because apparently more agents cooperate when using hierarchical initial-

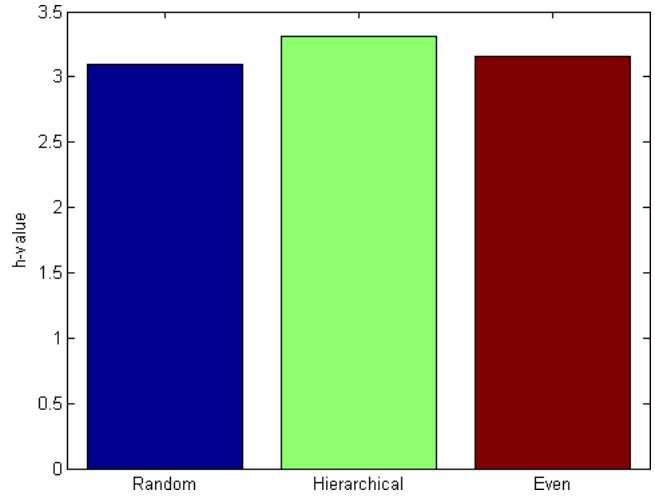


Figure 5: The mean hierarchical value of all agents.

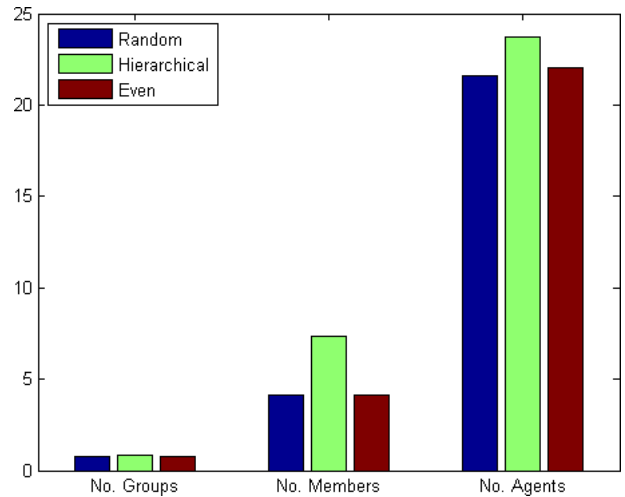


Figure 6: The mean no of groups, no of members of these groups, and agents in the simulation

isation (in bigger groups, not in more groups, as Figure 6 shows us), less agents disappear in the simulation, leading the overall number of agents in the simulation to be higher as well.

The influence of the initial threshold on the segmentation of the agents

The hypotheses here was that the initial cooperation threshold would positively affect the segmentation of the system (more groups with less members). According to Figure 7a, it clearly does. But concerning the number of groups, only till a certain point. We expected the amount of groups to increase when the threshold increased. This seems to be the case (when looking at the normal agents' threshold axis) till it reaches 1.25. After that point, the amount of groups drops, till nearly zero when the thresholds are 2.5. This is shown better in Figure 8. This is explainable. When the threshold is low, there is one big group with almost all agents as its members. When the threshold increases, agents will be pickier, and will cooperate less; there is not one big

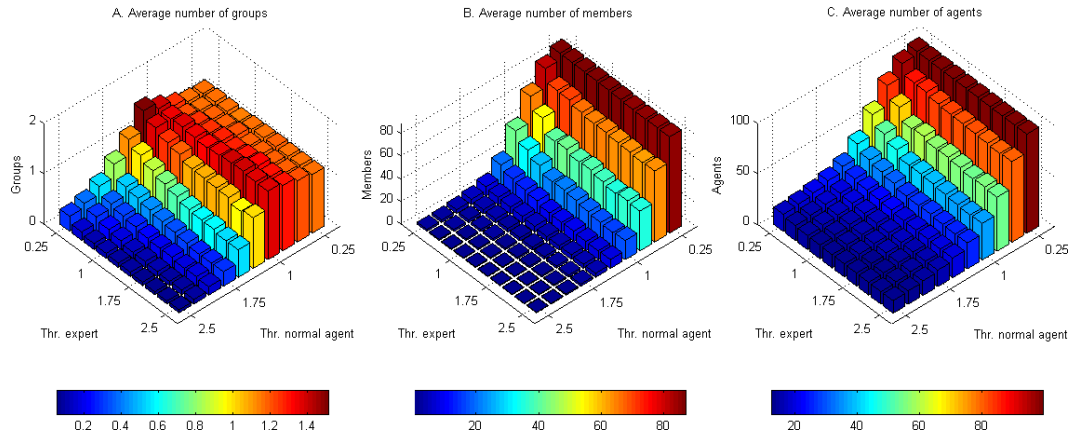


Figure 7: The mean no of groups, no of members of these groups, and agents in the simulation depending on the thresholds of both expert and normal agents. The different initialisations are averaged in the figures, because they were very similar.




Figure 8: The data of Figure 7a simplified. The data is averaged for the experts and the normal agents, so to show the influence on the h-value.

group, but several smaller. When the threshold increases even further, the agents do not cooperate at all anymore, or only in pairs or triangles (the group definition is 4 or more members). This causes the number of groups to plumber. This critical point seems to lay around 1.25 (for the normal agents’ threshold). Figure 7b confirms this. With a low threshold, the groups have a average amount of 80 members. Obviously this is one group, because of the maximum of 100 agents. The average amount of members decreases when the threshold increases. At the point where Figure 7a shows the most groups (threshold norm. agent =1.25) the average amount of members is around 10. After this point the amount of members reaches zero, just as the explanation of the number of groups predicted.

So the first part of our hypotheses (more groups) is correct only till a certain point, the second part (less members) is definitely true, and the last part (less agents) is confirmed bij Figure 7c. This figure clearly shows that the higher the treshold, the less agents in the simulation. This means that agents dissappear a lot, due to a lack of cooperations.

Important to noticed is that the influence of the normal agent is much higher than of the expert on the network, in

the number of agents, the amount of groups and the number of members, because the value seems to change much more in the direction of the axes of normal agents’ threshold than in the direction of the expert threshold. Figure 8 shows this as well.

When looking at the hierachal value of the network (Figure 9), we can conclude that a threshold of 1 or 1.25 is the most optimal for normal agents; it results in the highest hierachal value. This is something we found before when testing the model, and that is why the default setting is 1.25 for normal agents. The threshold of the expert does not seem to matter much.

What we were curious of, but cannot conclude is that there is a linear relation between the thresholds and the hierachal value. Apparently the hierachality of the network is not linearly dependent on the threshold. There does seem to be an optimal setting, which is in itself interesting, but does not contribute to this hypothesis and experiment.

The influence of the amount of experts on the segmentation of the agents

We expected to see a similar effect as in the previous paragraph when increasing the amount of experts; the percentage of the agents which have expertise. This time, the effect is less visible in the amount of groups (Figure 10a). Interestingly, the difference between the three initialisations is quite big here. The even initialisation has a slight higher group average, something we expected in the first hypotheses, but didn’t see there. On the other hand, the groups present when using the hierachal initialisation have more members.

We only see a slight effect of the amount of experts on the groups, and this effect is negative, instead of the expected positive effect. The slight grow of groups, and the decline after some point as seen in the previous experiment, is less visible here. An explanation can be, that because only a part of the agents have high thresholds, the cooperations mostly occur between normal agents, and experts dissappear quickly in the simulation. This would explain the drop in agents in the simulation when the percentage experts in-

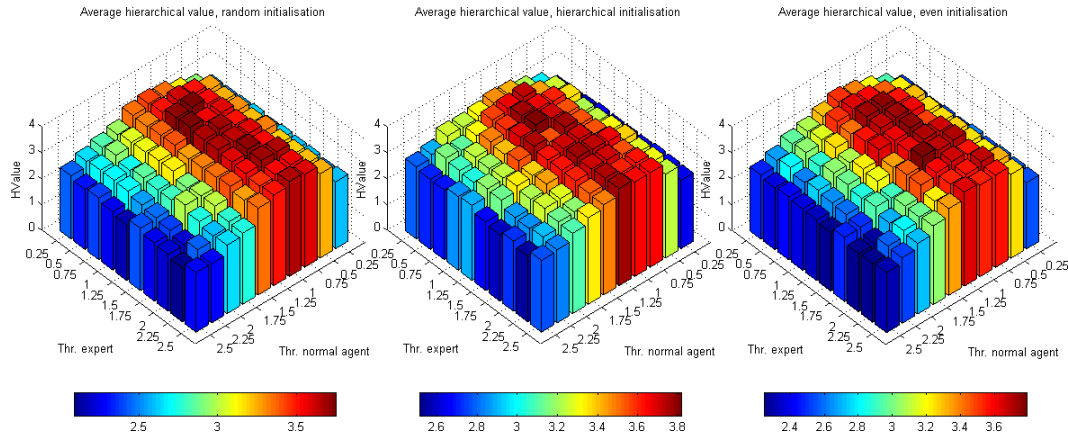


Figure 9: The average hierarchical value and the thresholds of both experts and normal agents.

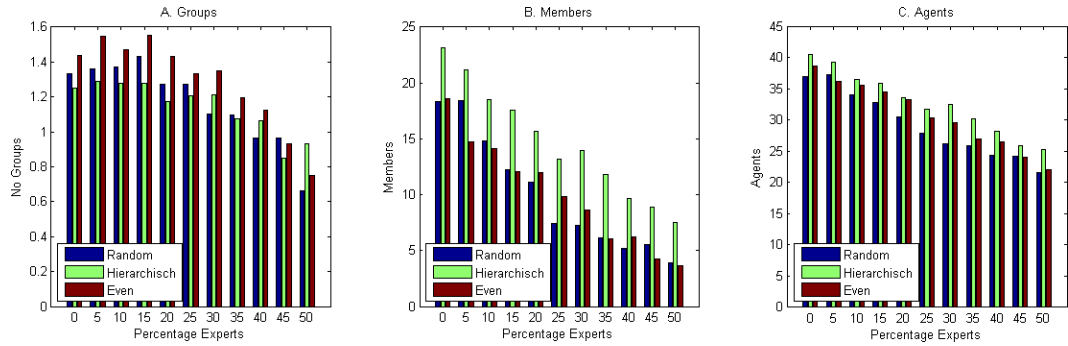


Figure 10: The mean no of groups, no of members of these groups, and agents in the simulation depending on the percentage experts.

creases, as is shown in Figure 10c, and would explain the amount of groups; less agents mean less groups. The effect of higher threshold as seen in the previous experiment apparently does not compensate for this tendency.

The reason that the hierarchical initialisation has less groups but more members can be explained by the fact that all agents with high value know each other (this is the essence of the hierarchical initialisation), while when using even or random initialisation, they might not. The chance is therefore higher that they will form one group, instead of multiple unconnected groups. This would also mean that the hierarchical initialisation will result in a higher average of colleagues per agent (one group has more connections than two separate groups of the same number of agents), and indeed, this is shown in Figure 11, where agents in the simulation with hierarchical simulation have a significant higher mean of colleagues (averaged for the percentage expertise); 3.47 colleagues per agent, compared to 2.56 for both random and even initialisation (two-sample $t(438)=6.4$, $p < 0.001$ in comparison with even initialisation and two-sample $t(438)=6.2$, $p < 0.001$ for random initialisation).

INTERPRETATION OF THE RESULTS

The conclusion we can draw from these three experiments is that the model does not provide evidence for the fact that the hierarchality of organized crime depends on how it started (the initialisation in the model). We expected the contrary, looking at for example Italian and American

mafia, where there was a clear opportunity which allowed the mafia to flourish; lack of government protection and trust in the government in Italy, and the prohibition in the United States.

The model does not give an answer to the question: “What determines the level of hierarchy of criminal organizations?”. However, the model does shed some light on the other question “What makes networks more flexible and fragmented?”. The results we saw suggest that this may be due to a higher tendency to distrust, which causes a high threshold for cooperations. In other words, maybe criminals in countries where we see the flexible kind of organized crime are more picky on who to cooperate with. Maybe the world of organized crime there is harder; cooperations are broken quicker, criminals do not cooperate with everyone, groups are therefore smaller, and criminals have a shorter stay in the criminal scene (due to arrests or other factors). It may point at the fact that for a hierarchical organization to occur the conditions have to be optimal: not too much competition, low levels of betrayal (high amounts of fear?) and/or no adequate police countermeasures.

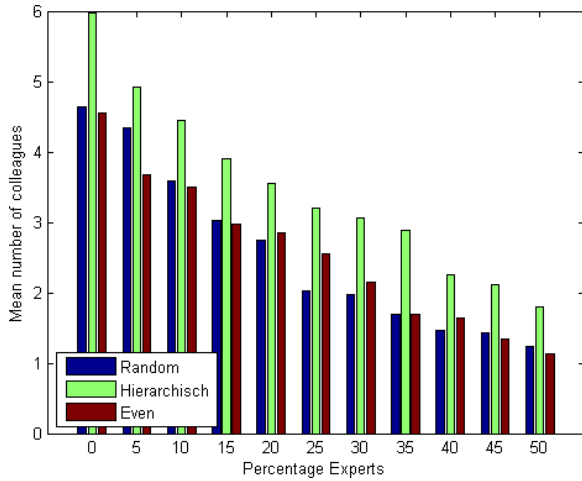


Figure 11: The average number of colleagues per agent, shown for all three initialisations in the percentage expertise experiment.

DISCUSSION & EVALUATION

Conclusions drawn on this model should be made with precaution. The model does not act all the way we hoped, and only replicates part of the phenomena we see in organized crime in the real world. This may be due to the following simplifications or choices.

One important aspect we ignore, is the environment in which the criminals cooperate. We do not look at the business model they use, the flow of goods, the type of crimes. Organized crime often uses the same opportunity structure that facilitates legal economic activities. Italian mafia is more known for its protection, because apparently the government can not provide this, while Dutch mafia specialises in international illegal trade; like drugtrafficking, human trafficking for sexual exploitation and other transnational illegal activities [2], using the prominent position of the Netherlands in the trading market.

The second weakness could be the initialisations. The reason why we chose for three different initialisation possibilities is because we wanted to rule out the influence of the initialisation on the resulting network. But one can also consider this the core of the simulation; is the difference between countries only due to its beginning? A topic to discuss now is if we can just only consider this ‘beginning of organized crime’ to be important at the start of the simulation. Maybe the beginning is not one moment, and this difference in how agents come into the simulation should be extended for a longer period of time. Now, new agents, when introduced during the simulation, are initialised the same for every initialisation; they get some random cooperations. Maybe we should introduce a hierarchical and flat way to add new agents during the simulation as well. Let an agent with high value connect with people of similar fitness when being introduced in the simulation, or, more likely, introduce only agents with relatively low value, and connect them to the bottom of the hierarchical network, so he can work his way up, but not find a place at the top right away.

Another point for discussion is the comparability of the initialisations. All three initialisation result in an average of two acquaintances per agent when the simulation is started. After ten time steps however, this quickly diverges into a more unequal number. Because hierarchical agents know people of around their fitness (value), the chance they will cooperate in the first epochs is bigger, except for the agents which are at the bottom (the leaves) of the hierarchical network. With random initialisation, some very fit agents may not have any connections, while unfit agents have many, resulting in a higher disappearance rate in the first number of epochs. The even initialisation will be somewhere between these, as each agents does have connections, but are connected with agents of arbitrary fitness. Maybe it would be better to pick a moment a little further in the simulation (say after 20 epochs), and adjust the initial relations in such a way, that at this point the amount of agents and cooperations in the simulation is the same. This may lead to more unbiased networks.

We chose to finetune some parameters beforehand, not changing them anymore during the experiments. We are talking about the snitching chance, the punishment of snitching, and the growfactor of reliability and capacity. In future research, these may be tested also to see their influence on the network, because they certainly have potential to influence the network in a meaningful way. The punishment and snitching chance can be varied for example as factors on how hard police is hunting these criminals.

Lastly, the importance of experts is still not as present as we hoped. This may be due to the the simple difference we gave them: a higher threshold. Experts are essential for criminal groups, so there should be a tendency to work with them when possible, but this is not the case in this model. An improvement can be that every group needs one to be ‘successful’. But how do you do this? Do you increase the snitching chance of a group (the police will find amateurs quicker than groups with experts) or do you implement another way in which a group is less viable when not cooperating with at least one expert? Here the flow of goods and succes argument comes back. When this is implemented you can let the fact if the groups contains an expert influence the profits they make or the success they have. This way, there will be a tendency to try to work with experts (more than with other agents), making their higher threshold justifiable; they are popular, so they can be picky who to work with. An argument against this is: how does an agent know if his group has an expert? He only knows his colleagues and their colleagues, but cannot see a level further. The top down behaviour you need to implement for this may stroke with the MAS notion of local interactions.

In this article we saw an example of the combination of criminology and multi agent systems. Although the research showed that there is still a long way to go in improving this cooperation for meaningful purposes, it did show the potential. The application of the useful tools artificial intelligence gives us on the field of criminology could bring new insights to well established theories, and by doing that, begin a new chapter in criminological research.

ACKNOWLEDGMENTS

The authors would like to thank prof. dr. P.J.V. van Calster for his advice and encouragement.

REFERENCES

- [1] A. Block. The snowman cometh: Coke in progressive new york. *Criminology*, 17(1):75–99, 1979.
- [2] H. v. d. Bunt and E. Kleemans. *Georganiseerde criminaliteit in Nederland. Derde rapportage op basis van de Monitor Georganiseerde Criminaliteit*. Number 252. Boom Juridische Uitgevers (Wetenschappelijk Onderzoek- en Documentatiecentrum): Den Haag, 2007.
- [3] D. Cressey. *Theft of the nation: The structure and operations of organized crime in America*. Harper & Row New York, 1969.
- [4] J. M. Epstein. *Generative Social Science: Studies in Agent-Based Computational Modeling (Princeton Studies in Complexity)*. Princeton University Press, 2006.
- [5] T. Fruchterman and E. Reingold. Graph drawing by force-directed placement. *Software-practice and experience*, 21(1):1129–1164, 1991.
- [6] C. Gerritsen. *Caught in the Act: Investigating Crime by Agent-Based Simulation*. PhD thesis, VU University, Amsterdam, 2010.
- [7] N. Gilbert. *Agent-Based Models (Quantitative Applications in the Social Sciences)*. Sage Publications, Inc, 2007.
- [8] R. E. Marks. Analysis and synthesis: multi-agent systems in the social sciences. *The knowledge engineering review*, 27:123–136, 2012.
- [9] E. Muller, J. van de Leun, L. Moerings, and P. van Calster. *Criminaliteit en criminaliteitsbestrijding in Nederland*. Kluwer, 2010.
- [10] X. Pan, C. S. Han, K. Dauber, and K. H. Law. A multi-agent based framework for the simulation of human and social behaviors during emergency evacuations. *AI & Society*, 22(2):113–132, 2007.
- [11] L. Paoli. De politiek-criminele nexus in italië: 150 jaar betrekkingen tussen mafia en politiek. *Justitiële verkenningen*, 35(3), 2009.
- [12] J. P. Scott. *Social Network Analysis: A Handbook*. Sage Publications Ltd, 2000.
- [13] P. van Calster. Netwerkonderzoek als perspectief op georganiseerde criminaliteit. *Justitiële verkenningen*, 34(5), 2008.
- [14] T. van der Beken. Kwetsbaarheid voor georganiseerde criminaliteit; een voor preventie bruikbaar concept? *Justitiële verkenningen*, 37(2), 2011.
- [15] R. van der Hulst. Sociale netwerkanalyse en de bestrijding van criminaliteit en terrorisme. *Justitiële verkenningen*, 34(5), 2008.
- [16] P. Williams. Transnational criminal organizations: Strategic alliances. *The Washington Quarterly*, 18(1):57–72, 1995.
- [17] M. Wooldridge. *An Introduction to MultiAgent Systems*. Wiley, 2009.