BILEL BENBOUZID, "Values and Consequences in Predictive Machine Evaluation. A Sociology of Predictive Policing", *Science & Technology Studies*, Vol. 31, 2018 (second version submitted after a first version accepted with minor revision). https://easst.net/sts_journal/

Abstract: *Predictive policing is a research field whose principal aim is to develop machines for predicting crimes, drawing on machine learning algorithms and the growing availability of a diversity of data. This paper deals with the case of the algorithm of PredPol, the best-known startup in predictive policing. The mathematicians behind it took their inspiration from an algorithm created by a French seismologist, a professor in earth sciences at the University of Savoie. As the source code of the PredPol platform is kept inaccessible as a trade secret, the author contacted the seismologist directly in order to try to understand the predictions of the company's algorithm. Using the same method of calculation on the same data, the seismologist arrived at a different, more cautious interpretation of the algorithm's capacity to predict crime. How were these predictive analyses formed on the two sides of the Atlantic? How do predictive algorithms come to exist differently in these different contexts? How and why is it that predictive machines can foretell a crime that is yet to be committed in a California laboratory, and yet no longer work in another laboratory in Chambéry? In answering these questions, I found that machine learning researchers have a moral vision of their own activity that can be understood by analyzing the values and material consequences involved in the evaluation tests that are used to create the predictions.*

Keywords: predictive policing; machine learning; algorithm; crime prediction

Predictive policing is a field of research whose principal aim is to develop machines for predicting crimes, drawing on machine learning algorithms and the growing availability of a diversity of data (Perry 2013). In the United States, predictive policing is part of a longstanding project of policing reform by research (Walker 2004) that seeks to create a pro-active police force which acts on its own initiative to prevent crime, rather than simply reacting in emergencies when called to do so by citizens. Since the 1970s, this reform project has in large part been driven by researchers critical of a situation where the police act mainly in contexts of crisis and drama, and are detached from concerns with preventing delinquency (Weisburd and Braga 2006). In this context, over the last four decades, a large proportion of policing research budgets has been devoted to experimentation on tactics that might allow the police to anticipate and precede the commission of offenses, rather than simply reacting to them. In 2012, when PredPol, Inc. put a predictive analysis platform on the market in the form of downloadable software, offering a dashboard that displays risks of crime in real time at a precision on the order of 200 meters (figure 1), the dream of the American police reformers of the 1970s seemed to have been given concrete form as a machine.

With the deployment of this type of analytical platform in public action, a new form of quantification began to progressively spread through public administrations. Classical public statistics, based on the law of large numbers and the associated notions of norms and means, now had to compete with these algorithmic practices, whose main objective is prediction through the automated production of classes, clusters, or patterns. Statistical learning was now liberated from the need for a fixed system of categories: "Rather than stable, permanent, structuring variables, which fixed statistical objects within categories, digital algorithms prefer to capture events that they record on the fly in order to compare them to other events, without first categorizing them. Instead of weighty variables, they seek to measure signals, behaviors, actions, performances" (Cardon 2016 p. 49). In recent years, as these changes have upset traditional reference frameworks in standard statistics, many initiatives have been undertaken to make algorithms a specific research object in the social sciences (Dourish 2016). In this literature, algorithms are seen as powerful mechanisms, with a growing role in all sectors of society and a subtle, discreet, and dissimulated power over individuals (Beer 2009). Denouncing

their intrusive, discriminatory, and underhanded nature (O'Neil 2016), researchers and activists have demanded a politics of algorithms (Crawford). Accountability (Diakopoulos 2014), transparency (Zarsky 2016), and audit (Sandvig et al. 2014) have become watchwords in these public debates on the algorithm.

The Foucauldian analyses of Rouvroy and Berns (2013) amplify what Ziewitz called an "algorithmic drama" (Ziewitz 2016). Rouvroy and Berns criticize the profound transformations in the exercise of power enabled by machine learning. They argue that the normativity of the law, in its discursive and explicit form, allows individuals the choice to obey or disobey and offers them the right to a fair trial which extends the possibility of dialogue; whereas machine learning imposes a "tyranny of the real" that neutralizes critique by producing normative devices based on strict descriptions of individuals' activity, or at least of their relations with their social and material environment. In other words, through the process of statistical learning, the "social norm emerges from the real itself". This algorithmic governmentality is characterized by its capacity to make all forms of resistance schizophrenic: discrimination, exclusion, and the unethical distribution of visibility are not directly produced by the classifications of the algorithms, but by the social reality on the basis of which the algorithms take form. Quantification is no longer the operation of institution of reality and transformation of the world that sociologists have sought to reveal (Desrosières 2002; Hacking 1990; Porter 1996), but instead an operation of conservation and reinforcement of that reality and of the flagrant injustices associated to it. In this context, actors' critical sense focuses on this algorithmic reinforcement of existing realities: that is, on the feedback effects of the computation, and not on the forms of computation themselves, with respect to which, according to Berns and Rouvroy, they lack all critical sense.

If actors lack any critical sense in these algorithmic contexts, neither the pragmatist sociology of social critique (Boltanski and Thévenot 2006) nor a sociology of controversies (Latour 1987) can be applied to machine learning algorithms. Does this mean the project of a sociological analysis of machine learning should be abandoned? Pointed as it is, Berns and Rouvroy's analysis in terms of critical dispossession limits the concrete possibilities for emancipation in relationship to algorithms of government, by confining actors, and sociologists themselves to a stance of powerlessness. If the sociology of science and technology is to contribute to the study of algorithmic prediction, and at the same time to justify its relevance and usefulness in this context, it must develop a specific art of inquiry that allows it to create critical tests specially designed for the purpose. Here, I borrow this art of inquiry from Tim Ingold (2013). To understand algorithms and their predictions, we must enter into a process of correspondence with them—touching them, manipulating them, and subjecting them to various operations.

I undertook such an inquiry between June 2013 and March 2017 in a study on the algorithm of PredPol, the best-known startup in predictive policing. The team of California researchers¹ behind it (Mohler et al. 2011) took their inspiration from an algorithm created by a French seismologist, David Marsan, a professor in earth sciences at the University of Savoie². As the source code of the PredPol platform is kept inaccessible as a trade secret, I contacted Marsan directly in order to try to understand the predictions of the company's algorithm³. Marsan tested his algorithm on the same open-access crime data from the city of Chicago that the California researchers used in their own publication. Using the same method of calculation on the same data, he arrived at a different, more cautious interpretation of the algorithm's capacity to predict crime⁴. Unexpectedly, I created a situation of controversy concerning knowledge of the technical properties of the algorithm (Mackenzie 2004). By confronting a physicist specialized in earth sciences with researchers in applied mathematics who are focused on developing predictive machines, I created an opportunity to take the beings who sustain the algorithm's existence and make them visible, and to focus my full attention on the specific associations that PredPol's algorithm is composed of (Latour and Venn 2002).

How were these predictive analyses formed on the two sides of the Atlantic? How do predictive algorithms come to exist differently in these different contexts? How and why is it that predictive machines can foretell a crime that is yet to be committed in a California

laboratory, and yet no longer work in another laboratory in Chambéry? To answer these questions is to describe a controversy that will allow us to explore the workings of the algorithm from the inside—to "unfold" it, in a Deleuzian term (Deleuze 2006) that Latour (ibid.) adopted for use specifically with technologies⁵. This procedure of unfolding revealed that machine learning researchers have a moral vision of their own activity (Daston 1995) that can be understood by analyzing the *values* and *material consequences* involved in the evaluation tests that are used to create the predictions. To analyze the moral dimensions of prediction is not to study this or that usage of machine learning, but to investigate the transformations undergone by predictive categories as they move from one social context to another. By the end of this article, prediction should be clearly understandable as a moral problem that is indissociably at once cognitive and material.

Figure 1: Screenshot of the map on the dashboard of the PredPol platform, indicating upcoming crimes at a precision of 200 metres.

AN ALGORITHM PREDICTING EARTHQUAKES AND CRIMES WITHOUT A PRIORI HYPOTHESES

What makes this California startup a reference in the domain of predictive policing is its slogan, "More Than A Hotspot Tool."⁶ Since the early 1990s, the urban cartography of "hotspots"—heatmaps of the distribution of crime in the city—has been the main tool for strategic intervention in areas where crime is concentrated (Weisburd, Bernasco, and Bruinsma 2009). PredPol claims to do better than these classical crime maps thanks to a predictive method used in the field of earthquake prediction. A similarity observed by these Los Angeles researchers between the dynamics of the propagation of crimes and that of earthquakes, they say, means that the geographical dynamics of criminality can finally be characterized mathematically (Mohler et al. 2011). PredPol uses a method of calculation taken from stochastic point processes, a branch of statistical physics. This is a classical approach to modelling the distribution of a set of events (considered as pointlike entities) in a finite space of arbitrary dimensionality (in the models discussed here, two spatial dimensions and one temporal dimension). Point processes are used to identify the formal mechanisms that produce these events in several dimensions, by modelling how they are distributed in time and space.

PRODUCING SPATIOTEMPORAL CLUSTERS BY COMBINING CONCENTRATION AND CONTAGION

The question behind this statistical operation is whether the events are distributed randomly or in some more regular fashion, and in particular whether the points cluster around particular locations. The choice of the type of process depends on the researcher's hypotheses regarding the form of mechanism involved. The PredPol researchers started with a classical hypothesis in predictive crime analysis: rather than occurring randomly, crimes are concentrated in space and spread through a local neighbourhood. The repetitive structure of the events themselves is enough to model them (without drawing on external variables). In other words, the best predictor of a future crime is a past crime. I will return to this hypothesis below, but note first of all that the PredPol researchers were interested in self-exciting point processes because they represent a way of modeling interactions between events that takes the history of previous events into account: the occurrence of future events (crimes or earthquakes) depends on the history of the process. The formula below is the mathematical representation of this process, the predictive algorithm used by Marsan (Marsan and Lengliné 2008) and PredPol (Mohler et al. 2011). It calculates a probability that, in a manner of speaking, represents an idealization of the mean number of crimes or earthquakes on a surface:

$$\lambda(x, y, t) = \mu(x, y) + \sum_{i, t_i < t} g(x - x_i, y - y_i, t - t_i, M_i)$$

In this formula, the probability of the occurrence of an event, at a given moment in the process, contributes to an overall calculation of risk intensity per unit surface and per unit time. In the language of statisticians, it is said that this equation “describes an inhomogeneous Poisson process of intensity $\lambda(x, y, t)$.” This intensity is a calculation of risk interpreted as a density that depends on both space and time. It is obtained by taking the sum of the two components of the formula: first, the function $\mu(x, y)$, known as the background rate, depends only on space, and represents a probabilistic calculation of the spatial concentration of risk in general; second, the function $g(x - x_i, y - y_i, t - t_i, M_i)$, known as the contagion kernel, models the spread of series of events whose occurrence depends on previous events and on the parameter M_i (the magnitude of the event). The algorithm models risk intensity at each location in a map by adding together these two components. Note that the contagion model is linear: if this were not the case, computation (simulation, optimization) would be very difficult, and would make the algorithm practically unusable. It is hypothesized, for example, that the contagion $g(1,2)$ following two distinct events 1 and 2 is simply the sum of the individual contagion effects of g_1 and g_2 (hence the sum in the equation).

A MODEL THAT ADJUSTS ITS OWN FORM

Marsan uses this method to model how the main seismic events in an earthquake set off aftershocks, which in turn set off their own sequences of earthquakes.⁷ The Los Angeles criminologists' interest in Marsan's algorithm was spurred by what they saw as a similarity of form with criminologists' characterization of crime dynamics: just as earthquakes are followed by aftershocks, so, they thought, crimes are followed by “aftercrimes.” Since the 1990s, a field of research on the repetition of crimes has developed in the United States and Great Britain, after a multitude of criminological analyses converged on the conclusion that most crimes repeatedly target a small number of victims, and propagate through their immediate spatial neighbourhood. In the model of “repeat and near-repeat crime concentration” (Pease and Tseloni 2014) that was proposed to account for this pattern, crime can be seen either as the signal of a relatively stable risk in a given area, or as an indication that incidents of victimization reinforce the probability of

the occurrence of later incidents: In other words, crimes are repeated in or near the same location, and spread by “contagion.”⁸ These two hypotheses are present in the calculation of risk intensity presented above: the relatively stable risk in an area corresponds to the concentration $\mu(x, y)$, and local reinforcement to contagion $g(x - x_i, y - y_i, t - t_i, M_i)$.

With this algorithm, Marsan seeks to show that the structure of cascades of events can be modeled probabilistically, without any particular hypothesis about the underlying mechanisms, and without the need to test model parameters first. Herein lies the contribution of Marsan’s algorithm in seismology. While most existing seismological models are parametric (see below for an explanation of this term), with the parameters set on the basis of empirical data, Marsan and his collaborator made their mark in the field by showing that this parameterization can be dispensed with completely. The PredPol researchers made the same argument to justify the value of their own research: the first statistical approaches used in criminology, notably in the study of the spread of crimes, were parametric. This means that they required a hypothesis on how crimes propagate. But how exactly do parametric and non-parametric approaches differ?

In parametric models the form of the model is imposed and its parameters optimized. In nonparametric models, in contrast, an estimate is calculated of the optimal size of the diameter of the circular moving window (smoothing window) that records the number of points in each cell in a virtual grid projected on the map, and the number of parameters varies, increasing with the number of observations. To estimate the parameters (“nonparametric” does not mean free of parameters), Marsan uses the expectation-maximization algorithm, a classical method that, in an iterative procedure, repeatedly alternates two steps (the calculation of expectation and the calculation of maximum likelihood), in order to arrive at the estimator of the model. Marsan refers to another better-known method, artificial neural networks,⁹ to explain how this non-parametric method follows in the spirit of machine learning (Domingos 2017) :

It's a little like a neural network. We put in bricks that depend on parameters, but the final product is not, or is very little constrained at the outset. The model adjusts its own form. For us, it was above all a way to show that a model with the fewest possible assumptions could converge toward laws (forms) that are very close to the empirical laws conventionally injected at the beginning in stochastic approaches. The fundamental difference with neural network approaches is that they're often used as a black box with a strictly predictive goal, whereas that's not at all what we had in mind. Instead we're trying to understand what a "good" contagion kernel is and how it can emerge naturally from the data analysis.

According to Marsan, this opposition between understanding to construct theories and predicting to act without necessarily having a complete understanding of the phenomenon—which is well known to machine learning specialists (Shmueli 2010; Hofman, Sharma, and Watts 2017)—explains the difference between his approach to modelling and that of the PredPol mathematicians. Here we will look at the practical consequences of this opposition for the way the algorithm is evaluated.

REVEALING THE SUMPTUOUS OPACITY OF THE ALGORITHM

Through this simple surface description of the algorithm, we have surpassed the barrier of its mathematical formalism and glimpsed at the hypotheses that its predictions depend on. Here I will move on to a sociology of the knowledge of the technical properties (Mackenzie 2004) of PredPol’s algorithm. To do so, I will consider the algorithm as a being whose dominant mode of existence is technological (Latour 2012; Simondon 2016). As Latour clearly showed, “The technological object is opaque, and—to put it bluntly—incomprehensible [...] in that it can only be understood provided that we add to it the invisibles that make it exist in the first place,

and that then maintain, sustain, and sometimes neglect and abandon it." (Latour 2010 p.26). From this perspective, it can be clearly seen that opacity, which is now a commonplace in public debate on algorithms, is not a problem specific to machine learning: all technological beings, generally speaking, "like to hide." There is no use hoping that the PredPol algorithm will become transparent, that its developers will make it public in order to clarify and make it easier to master: like all technical beings, algorithms are fundamentally opaque. Nevertheless, an appropriate method of inquiry can allow us to get into the workings of the algorithm, revealing its "sumptuous opacity" (Latour *ibid*).¹⁰

TOUCHING THE ALGORITHM

To unfold the PredPol algorithm, we must know how to use its language—not the specific language of code, but the more general language of *the technical*: detours, zigzags of ingenuity, ruses (Latour 2012). Faced with the opacity of the PredPol algorithm, we reacted in something like the way an archaeologist might when faced with an ancient object whose meaning escaped her: turning it around to view it from various angles, simulating it, and reproducing it. This stance is an unfamiliar one for the sociologist: handling the object of inquiry, squeezing it, fiddling around with it, hacking it. I thus asked David Marsan to do this on my behalf, and discussed it with him on several occasions. He tested the algorithm that I have just presented on open-access crime data from the city of Chicago—which, as noted above, is the same data that the PredPol researchers used in one of their own publications. I thus asked Marsan not only to explain to me how the algorithm works, but to run it on crime data and share the results with me. Unexpectedly, in so doing, I created a situation of controversy around knowledge of the technical properties of the algorithm. Nothing like a good controversy to get inside the workings of a machine.

To describe this controversy, I will follow the three steps that structure the work of modelling itself: the justification of the choice of point process (here, self-exciting), and then the modelling strategy and the associated model estimation techniques (the expectation-maximization algorithm), and finally, the evaluation of the model. The first two steps depend on analysts' beliefs about the nature of the problem, whereas the last depends on their moral vision of their own activity. In the experiment that I proposed to Marsan, the type of point process was imposed, and his modelling strategy did not significantly differ from that of the PredPol researchers. Marsan thus set out to evaluate a model applied to a phenomenon about which he knows nothing, using an algorithm similar to PredPol. Nevertheless, points of divergence appeared when it came time to interpret the results. Marsan expressed numerous doubts on the capacity of his algorithm to do better than classical maps of crime hotspots. In the note that he wrote on the analysis of the Chicago crime data, he concluded:

The results obtained offer good reasons to doubt the capacity of the proposed models to do better than simple hotspot maps. The contribution of contagion (the triggering contribution) to explaining the occurrence of future events is small (it represents only 1.7% in the best model). The role of "memory" in the process can thus make no more than very modest contribution to the efficiency of the prediction system. More importantly still, the assumption is that the dynamic of the process remains the same over time. The possible non-stationarity of the process is clearly a problem, because it limits the use of past information to predict the future. In 2015, burglaries were not distributed (in time and space) in the same way as in 2014. This non-stationarity is probably due to uncontrolled changes in how criminal acts are carried out. It could also be due to the deployment of new predictive algorithms: as police patrols use them, they might provoke reactions among burglars. Contrary to natural processes such as earthquakes, analyses like the ones presented here could change the observed process, which makes correctly predicting future events more difficult (personal note from David Marsan, sent to George Mohler in September 2015).¹¹

To understand Marsan's conclusion, recall that the algorithm calculates the intensity of risk in space and time by adding together two elements: concentration (space) and contagion (space- and time-dependence). David Marsan's note indicates that contagion does make a contribution to the process, but it is extremely small—in fact, negligible. And yet this is the dimension emphasized by the promoters of PredPol in their slogan "More Than a Hotspot Tool." Could it be that the PredPol scientists altered the results to make them more favorable to their commercial project? According to Marsan, the answer is no: the PredPol researchers did honest work. Moreover, Marsan wrote to George Mohler, who responded as follows:

Thanks for your email and sending along the analysis. I have found your work on nonparametric point processes quite interesting and influential! We have certainly seen the branching ratio vary quite a lot from city to city and crime type to crime type (from 0 to .5). As you point out, it is important to pick such parameters using cross validation in which case it is certainly possible that a simpler model may be favored. It also may be the case that the nonparametric model you are using is over-parametrized (it looks like it has over 30 parameters), so it may be over-fitting the training data. You might need more regularization, or you might want to use a semi-parametric model (you mention using an exponential smoothing kernel, which is essentially a parametric Hawkes process without the background rate). Another thing you bring up is the non-stationarity of the process. I think this is important and something we tried to estimate in the JASA paper (Mohler et al. 2015) (where the background rate μ depends on time). Disentangling endogenous contagion from exogenous fluctuations in the intensity is a somewhat open problem, though I have done a little work in this area. The non-stationarity of the background rate is one big difference between crime and earthquakes, and you often try to factor in seasonality and other explicit exogenous predictors. (Email from George Mohler to David Marsan, 3 September 2015).

There is no reason to question the honesty of the scientists who worked to develop PredPol. In his response, Mohler shows he is conscious of the many limitations of the PredPol algorithm, and offers a defense against Marsan's critique, recalling that they sought to deal with the problem of non-stationarity by adding a time variable $v(t)$ to the background rate $\mu(x, y)$. The PredPol equation thus becomes $\lambda(x, y, t) = v(t) \mu(x, y) + \sum_{i, t_i < t} g(x - x_i, y - y_i, t - t_i, M_i)$. How did Marsan react to Mohler's response? To answer this question we have to follow in the steps of Marsan's critical analysis:

The little work I did on this - well, it took two or three weeks of work, that's not nothing - showed me that there was a problem in the data between 2014 and 2015. I took a look in a very simple way to see how they behaved, and in fact they aren't similar at all.

Marsan then invites us to look at the two figures below. On the left are earthquakes; on the right, crimes.

Figure 2. Graphs showing the evolution of the mean distance between pairs of events, from Marsan's notes. These graphs illustrate the "memory effect" and the stationarity problem. The x-axis is exactly the same on both graphs.

These two graphs, which represent simple descriptive statistics on change over time in the mean distance between pairs of events separated by $n-1$ events, led Marsan to say that the data do not behave in at all the same way between 2014 and 2015. According to Marsan, the memory effect is very weak for crime. While the distance between pairs of events increases with the number of interposed events for earthquakes, for crimes it is not apparent that there is any such trend. The most surprising thing for Marsan was that for crime, the mean distance was quite different for different years. The phenomenon is not stationary.

ACCURACY VS. PRECISION

How is it that these simple descriptive statistics were enough to convince Marsan that his algorithm is not particularly applicable to the Chicago data, but that they did not concern Mohler, who—as his response to Marsan shows—was not surprised by this difference between seismological and criminological phenomena? Marsan suggested a partial response in our interview:

Bilel, you have to understand. You're a statistician, you don't know much about the problem you're being given, and they say to you, "We'll pay you, we'll give you the data, give us the best possible model." You go to work and you realize that your model behaves well one year and then a year later it doesn't. You're a stats guy, you don't know much of anything about the problem. What do you do? He as a statistician says to himself, my model isn't flexible enough, I'm going to make it a little more flexible, I'm going to add $v(t)$. Personally, I'd rather go and talk to the Chicago police to try to understand what happened, what changed. Why is it different in 2015 than it was in 2014? Is it a counting problem? Did the police officers change their habits? Basically you try and understand what made it change from year to the next. Maybe Mohler tries to understand, but his attitude suggests to me that that's not what he does. He tries to improve the predictive power of his algorithm. But because that doesn't work too well, he tries to make it a little more flexible so it works better. His model isn't flexible enough, so he says "I'm going to take my $\mu(x,y)$ and make it a bit more flexible so it works better by adding temporal variation to the background rate".

Thanks to his intimate knowledge of the algorithm, Marsan was able, in a manner of speaking to get inside his own algorithm, putting himself in Mohler's skin. Before Marsan looked into the Chicago data, the PredPol algorithm remained invisible not because it was protected as a trade secret, but because everything that would make it possible to follow the algorithm's course of action remained hidden. Marsan allowed us to bring out some of the invisibles (Latour *ibid.*) that the algorithm depends on. In doing so, the seismologist revealed the ingenuity with which his research was diverted, transformed, and translated to become usable in a police officer's smartphone. He discovered, with the same stupefaction as me, what had become of his algorithm in the hands of a team of audacious mathematicians. Nevertheless, Marsan was critical of the way in which the PredPol developers deployed his machine:

It might be that that's not the right approach. It might be that it's even the contagion that's different from one year to the next. You'd have to switch out the contagion kernels. But that's the hardest part to adjust. It's simpler to just add a time variable. What he does is really basic. In seismology, we do things that are much more complex to get the background rate to change over time, to take non-stationarity into account. The essential step after the PredPol article would be to understand the non-

stationarity. But they're driving blind. Personally I think you can't analyze your data without asking questions about the reality they represent. If you like, the two of us aren't driven by the same engine. What interests us in seismology isn't doing prediction, it's understanding the form of the kernel. Contagion interests us because it gives us clues about the mechanisms that make it so that one earthquake sets off another. It interests us because it tells us something about the seismogenic process. We're not going to impose an a priori form, because the form is what interests us. He's not interested in the form of the contagion. He doesn't want to understand how the contagion happens. He wants to make a prediction. It's totally different. In our field you can find the same kind of researcher. There are people who do prediction, but who don't want to understand the process. And there are many of us who think this leads you into a dead end.

Here, Marsan is no longer discussing the effectiveness of the machine in relationship to the crime data, he is offering a *moral condemnation* of the work of the PredPol developers. In his view, he as an earth sciences professor in France does not share the same values as the generalist mathematicians in California. According to Marsan, prediction poses a "basic" research problem, in that what is in question are the theoretical foundations of seismology in a context where, as theoretical knowledge currently stands, predicting earthquakes is impossible. He repeated the point several times in our interviews:

Short-term predictions (from a few hours to a few days) are rarely successful. Most times they're cruelly disappointing. Our failures at prediction regularly raise the question of whether predicting earthquakes is fundamentally impossible.

Marsan thus sees basic research as a tactical retreat: the idea is not to claim to produce pure and autonomous research, sufficient unto itself, but to take a step to the side into more theoretical research in order to overcome the problem of earthquake prediction. The community of seismologists to which Marsan belongs argues that a comprehensive approach to the phenomenon is needed. They oppose exclusively probabilistic and predictive analyses, which are defended by researchers who consider that theoretical research in seismology has reached its limits (cf. the debate in *Nature* [Main 2017], and notably the opposing positions of Pascal Bernard and Didier Sornette, which Marsan highlighted in our interview). Marsan analysed the crime data in the spirit of his work on earthquakes. On his view, a nonparametric model is appropriate if the patterns yielded by this statistical learning approach raise important research questions. This principle implies an evaluation of the "accuracy" of the algorithm: what is evaluated is the capacity of the calculation to reveal a close link, or a certain degree of accuracy, in the match between the mathematical model and a coherent conception of the phenomenon under study. The PredPol researchers judge the algorithm according to different criteria: if it improves the "precision" of prediction scores, then the algorithm is satisfactory.¹² Marsan and the PredPol developers do not subject the algorithm to the same tests. In California, the crucial test is performed through a type of lift curve, which is a tool for comparing the performance of different algorithms (figure 2). It was in this spirit that the PredPol developers turned to point process statistics, testing the algorithm on a criterion of competition with other algorithms on the prediction market. These California mathematicians aspire to other principles, basing their research practice on objectives of precision, efficiency, simplicity, and the ability to bring a predictive solution to market—all values that may be totally detached from the accuracy and correctness of the results produced by the calculations (Daston *ibid.*).

Figure 3: An example of a curve of predictive efficiency used in the article of Mohler (2011), showing the quantity of successfully predicted crimes in relationship to the number of cells flagged. The graph demonstrates the predictive superiority of PredPol with respect to Promap, the pioneering predictive policing tool developed in England in the 1990s by the criminologist Ken Pease and his collaborators.

THE ROBUSTNESS OF A PREDICTION IS INVERSELY PROPORTIONAL TO ITS PRACTICAL CONSEQUENCES

One invisible that would have been difficult to detect without the help of David Marsan is the fact that the PredPol developers added flexibility to the algorithm by integrating a time variable to the background rate, thereby resolving the thorny problem of non-stationarity in the simplest possible way. Through this unfolding of the algorithm, Marsan freed us from the widespread obsession with the question of predictive efficiency alone. Algorithms are technological beings that can offer the opportunity for a much richer debate. The situation of controversy that I created from scratch (a useful way of demonstrating the sumptuous opacity of the algorithm) revealed two different moral visions of predictive activity: one focused on the correctness (or accuracy) of models, the other on the precision of risk scores. Now what is needed is to follow the lines of the network that is laid out on the basis of these two different ways of assigning value to predictions.

CONCEPTUALIZING THE PROCESS UNDERLYING A PHENOMENON OR CAPTURING THE LARGEST POSSIBLE PROPORTION OF EVENTS?

Recall this fundamental principle in the sociology of the sciences: phenomena are defined by the response they give to the tests that scientists subject them to in their laboratories (Latour, Woolgar and Salk, 1992). In the Chambéry earth sciences laboratory, what Marsan calls an "aftershock" acts as a specific being:

When a seismologist is analysing aftershocks, he doesn't content himself with counting them. First of all, periods of high activity are the ones when he has the best chance of catching a large earthquake in the net of his measurement networks. If there are enough recordings of good enough quality, he'll be able to establish a tomography of the rupture of the fault. Even without major aftershocks, he'll learn a lot from the small ones, particularly on the directions of tectonic constraints, which he can deduce from their mechanisms. The bulk of the analysis work consists in localizing earthquakes: based on the arrival times of the P waves and S waves from each earthquake at each seismometer in the network, you draw the cloud of aftershocks point by point. It looks

heterogeneous, but it includes calm areas and swarms, whose distribution changes over time: the subterranean cloud moves progressively further away from the epicentral zone. The mechanical disturbances of the main fault decrease quickly with distance, so that aftershocks are mainly only observed in its neighborhood, at distances equivalent to the length of the fault itself. By studying the aftershocks you can draw a circle around the main area of movement and localize the fault that's responsible, even if it doesn't break all the way to the surface. Even better: the form of the cloud that they draw out in the opacity of the earth can sometimes be used to determine the orientation of the fault. In certain cases, an abnormal concentration of aftershocks reveals the beginnings of the destabilization of neighboring faults. In the early 1990s, it seemed like everything had been done or said with aftershocks. Post-earthquake field studies became a matter of routine, a well-oiled machine, with ever more effective standard analyses. However, while the images of aftershocks were becoming more precise, detailed interpretation seemed to be impossible, as it depended on uncontrollable parameters linked to resistance and to the—unknowable—state of tension of peripheral faults." (Bernard 2003 p.108, reading suggested by Marsan.)

Seismologists see aftershocks as a chance to understand what happened, to "catch a large earthquake." In Marsan's more technical formulation, aftershocks serve to "filter the signal." In seismic catalogs, all seismic waves have been recorded jointly by the seismographs. This is why seismologists need to isolate independent earthquakes (for example, those linked to the secular movements of plate tectonics, also known as mainshocks from earthquakes that depend on one another (foreshocks, aftershocks, multiplets). Marsan's algorithm is one of a range of methods known as *declustering* methods, which were developed to try to capture independent earthquakes when analyzing catalog data, distinguishing them from all other seismic events, notably those that correspond to aftershocks. The main challenge is to enrich the catalog in order to model seismicity as a process, wherein the occurrence of one earthquake alters the surrounding tension field and the capacity of nearby faults to generate other earthquakes. To model seismicity as a continuum of earthquakes, Marsan must be able to isolate classes of earthquakes, in order to integrate the fact that the tensions released by small earthquakes may be as large as those resulting from larger earthquakes in the locations where the seismicity occurs. It is in this spirit of isolating classes of earthquakes that Marsan observes the contribution of "memory" to seismicity in an area. In a critical methodological article on declustering algorithms, Marsan and coauthors surveyed the advantages of statistical learning, which free seismologists from the need to define a priori the statistical characteristics of the classes of earthquakes that seismicity consists of. Marsan is interested in artificial learning because it allows him to challenge the system of categories that he uses to investigate seismicity. He approaches the existence of the three classes of earthquakes—foreshocks, mainshocks, and aftershocks—with a certain methodological nominalism. For him, declustering must be used to test whether the conventional forms of earthquake classification are well founded. He closes the article with the following lines on this epistemic opening:

Even though great progress has been made in the last decade, there are still many open questions, i.e., starting with the physical triggering of earthquakes (aftershocks), effects of uncertainties in the catalog on the results of declustering, or the effect of censored data (selection in time, space and magnitude range) on the outcome. In summary, care should be taken when interpreting results of declustering or results that depend on a declustered catalog, because these results cannot reflect the exact nature of foreshocks, mainshocks and aftershocks; indeed the exact nature of these events may not exist at all. (Van Stiphout, Zhuang, and Marsan 2012 p.24)

What if foreshocks, mainshocks, and aftershocks did not exist before being modelled? Marsan takes such an "agnostic" approach to modelling. He proposes to suspend beliefs regarding earthquakes, abandoning the idea of a pre-data structure that can simply be observed in the

catalogs. In other words, Marsan expects machine learning to be able to be placed *underneath* categorical forms of seismicity (Cardon 2016). At no point did Marsan see the algorithm that he programmed as a method of predictive analytics, because predicting aftershocks is not an end in itself in his research.¹³ Aftershocks interest Marsan because they have the power to help him conceptualize the process of seismicity differently.

In their applied mathematics laboratory in Los Angeles, the PredPol developers use aftershocks in a different way. In their article they suggest that declustering methods can offer a means of enriching "crime catalogs," but do not expand on the point. Using prediction as a declustering method, as Marsan does for earthquakes, could contribute to research on the modelling of crime in general—a subject I have written about elsewhere, and will not pursue further here (Benbouzid 2015)—which poses basic research problems no less complex than those of seismology. But the PredPol researchers are interested only in the "aftershocks" for the possibility they offer of adding an additional process to hotspot maps to incorporate regularities (repetitions). They use "crime aftershocks" (near repeat crime or near repeat victimization) for their capacity to capture the largest possible proportion of events.

THE VALUE OF A PREDICTION IS INSEPARABLE FROM ITS PRACTICAL CONSEQUENCES

Thus, on the spatiotemporal projection traversed by the algorithm, repetitions are what they do as a function of what scientists try to make them do. Between California and Chambéry, the status of repetition of crime changed, because they took up a place in two different institutional environments. In Chambéry, measured classes of entities exist in a domain where predictions lead to demonstrable consequences, which is not true in the case of policing. To illustrate this situation, in our interviews, Marsan often opened a historical parenthesis on the problem of demonstrable consequences in earthquake prediction. He mentioned the example of Parkfield as a symptomatic case of this problem in seismology:

Parkfield is a little village on the segment of the San Andreas Fault that seismologists have transformed into an observation site, which is now considered the most instrumented place on the planet. An earthquake was predicted there in 1988, but it happened in 2004, 16 years late. Failures of prediction of this kind are not rare in seismology. They pose particularly serious problems.

In the politics of earthquake prediction, scientists are held directly responsible for false positives and false negatives. The experience of false positives (earthquakes that do not happen) leaves inhabitants with a feeling of generalized anxiety and causes large economic losses. The experience of false negatives, as in L'Aquila, Italy in 2009, clearly explains seismologists' reserved attitudes on their own capacity to provide robust predictions. Predicting an earthquake implies evacuating entire cities, which carries a considerable cost and can provoke dangerous panic reactions at the scale of the road network of an urban agglomeration.

In crime prediction policies, the PredPol platform works as a tool for the management of police action. PredPol's research has shown that by spending just 5% of their available time in the areas identified by the algorithm, police patrols are twice as effective (in terms of crime reduction) as when they patrol the hotspots classically identified by analysts. The accuracy of PredPol's claims is not very important. What counts is to be able to optimize, and above all to *precisely* control, this tactical allocation of police time to presence in high-risk space. By integrating data from the GPS tracking systems installed in police vehicles, the algorithm optimizes the dosage of the presence of police patrols in different sectors of the city: the predictive square remains red on the map as long as the police have not patrolled there, turns blue during their first movements through the area, and then green when the officer has spent the optimal period of time as calculated according to available resources (for example, 5% of a police officer's working day). For a sector manager, PredPol appears to be a good tool to ensure that police officers play their preventive role, often simply by way of their dissuasive presence,

distributed randomly, but for an optimized duration, in the areas where risk is estimated to be greatest. The task of prediction is the management of the public supply of day-to-day police presence, while minimizing the need for change in police organization. While earthquake prediction has profound effects on the material and social structures of a city, crime prediction, as PredPol sees it, involves a minimal transformation in how policing is organized. David Marsan's meticulous, reserved, and prudent attitude can be understood as a habit developed in a field where researchers are held responsible for predictions that may have serious consequences. In contrast, the runaway success of PredPol can be attributed to the limited practical consequences of its predictions—hence the relatively casual manner in which the California researchers claim to predict crimes. This observation fits well with what sociologists have shown in studies of “theories, machines, and technology”: that “their robustness, their solidity, their truth, their efficiency, and their usefulness depend less on formal rules or on their own characteristics than on their local and historical *context*—this independently of the various ways that there are of defining that context (Teil and Latour 2017). The robustness of earthquake or crime prediction is not the result of a rational calculation, validated by neutral researchers and integrated into a machine. It is a solidity composed of the actions targeted by the prediction and the network of the material elements that they imply. The moral of this controversy is that the robustness of a prediction is inversely proportional to its practical consequences.

THE DIVINATORY ASPECT OF MACHINES (CONCLUSION)

To go further in this examination of relations between predictions and their consequences, a good source of inspiration is an article by the anthropologist Joel Robbins on deontological and consequentialist styles of reasoning (Robbins 2010). Robbins refers to this style of reasoning, based on “appropriate rules and not on the consequences of one’s rule-governed actions,” as deontological. He contrasts this approach to morality with consequentialism, where “actions are judged by their results, not by how closely they conform to a given rule.” Robbins then deepens this analysis of deontological reasoning, drawing on an article by Jane I. Guyer, “Prophecy and the Near Future,” on what she calls “the evaporation of the near future in theory and public representations.” Guyer shows how, in both contemporary economic policies and Evangelical discourse, the dual focus on the immediate present and the very long-term has taken the near future out of play as a temporal frame. Robbins adds to Guyer’s analysis that “[w]hat is lost in this move is the provision of a temporal space for [...] consequentialist reasoning,” in favour of “deontological forms that do not need to refer to the near future world of demonstrable consequences to reckon the value of actions.” According to Robbins, the success of evangelicalism can be attributed to the way in which the contemporary period creates uncertainty about the near future, which can no longer be predicted at all, leading individuals to concentrate on the present of their actions and to project themselves into a distant and mystical future. Here, respect for principles wins out over the anticipation of effects: “Different styles of moral reasoning are embedded in different kinds of social circumstances, and [...] forms of moral reasoning only flourish in those social circumstances that are well suited to them. Consequentialist moral reasoning, for example, only works where people have a sense that the social world they inhabit is relatively predictable, such that the probable consequences of an action appear relatively easy to gauge with certainty. Where such conditions do not hold, deontological approaches make much more sense—even in situations in which one cannot control the consequences of one’s actions, one can control whether or not they conform to a rule or set of rules (Robbins *ibid.* p.124).”

This distinction is of interest here, as it analyzes two different moral approaches in relationship to forms of prediction, in direct parallel to the differences between the approaches of Marsan and PredPol. The social circumstance of the unpredictability of earthquakes might seemingly favour deontological approaches, but in reality Marsan's stance is consequentialist:

Marsan pursues his research in the aim of making earthquakes more predictable, and thus to confer intelligibility on public announcements of the probabilistic theoretical construction of a phenomenon. PredPol gives the police the feeling of working in a more predictable world, but situations where police officers can directly observe criminals in the act are rare, even during discreet undercover patrols in the areas indicated with a precision of 200m x 200m. How, then, can PredPol claim to "predict" crime? A remark made by Sean Malinowski, the first Los Angeles Police Captain to experiment with the PredPol platform, offers a glimpse into what "prediction" means for the police: "If honestly done, there are no bad predictions in crime control."¹⁴ Contrary to the seismologist, police officers cannot experience "failed" predictions, because in their practice, prediction is expressed not in terms of truth or falsehood, but in terms of "good" and "bad." The problem is not to believe or disbelieve in the machine's predictions, but to do something rather than nothing, following the machine's recommendations.¹⁵

It could be argued in response that the PredPol researchers adopted a consequentialist ethic, as they implemented a system evaluating the effectiveness of the algorithm's recommendations (Mohler et al. 2015). But this evaluation bears only on very short-term consequences, and does not test the statistical significance of the measured decreases with respect to the general trends in crime over the long term.¹⁶ It is difficult, or even impossible, for the police to assess the practical effects of their daily activities on long-term trends in crime. The PredPol software allows them to optimize their attempts to control an overwhelming social phenomenon (Manning 2008). It is simpler for the police to rely on the dosages recommended by the machine: through the statistical learning procedure, crime takes form in a machine that produces rules to be followed by the police. PredPol remains a self-enclosed automaton, and like all automatons, it can provide only summary results (an analytical dashboard indicates upcoming risks by simply adding an additional process to hotspot maps to incorporate regularities (repetitions)). The seismologist, in contrast, makes use of the inductive logic of machine learning to take on a role as the continuous "regulator" and "organizer" of the predictive mechanism. In Marsan's hands, the algorithm became an *open machine* (Simondon 2012), whose functioning can be deliberately modified, and which is used to *understand*.

To summarize, what distinguishes Marsan's approach from that of the PredPol developers is that the seismologist conceives prediction in terms of its practical consequences, and the developers conceive it in terms of an absolute duty to act. Predictive policing is deontological insofar as the principal question that it asks of the algorithm is "What must I do?" and not "What is the best possible world with respect to the consequences of my actions?" (Ogien and Tappolet, 2009). These two moral approaches to prediction are applied in two different practical temporal spaces. Marsan conceives of prediction on the scale of the near future, the time needed to provide supplies to an area or to evacuate a city, a time frame that requires him to conceive the moral dimension of his research activity in terms of foreseeable consequences. His ethic of responsibility pushes him to say that "We're incapable of prediction." PredPol's predictions are focused on the immediate present—real-time analysis—and not on the long-term consequences of the actions that are organized by those predictions. The attitudes of the Evangelicals¹⁷ in Robbins's study and the integration of predictive machines into police organization may, in this way, be more similar than it seems. The former locate the future in the hands of God, the latter in the hands of a machine that police leaders hope will lead to salvation. When they work according to this deontological style of moral reasoning, predictive artificial learning machines are made not only of technology, science, and organization, but also of an element of divination.

BIBLIOGRAPHY

- Beer, David. 2009. 'Power through the Algorithm? Participatory Web Cultures and the Technological Unconscious'. *New Media & Society* 11 (6): 985–1002. <https://doi.org/10.1177/1461444809336551>.

- Benbouzid, Bilel. 2015. 'From Situational Crime Prevention to Predictive Policing'. Translated by Nonta Libbrecht-Carey. *Champ Pénal/Penal Field*, no. Vol. XII (March). <https://doi.org/10.4000/champpenal.9066>.
- . 2016. 'Who Benefits from the Crime?' <http://www.booksandideas.net/Who-Benefits-from-the-Crime.html>.
- Bernard, Pascal. 2003. *Qu'est-ce qui fait trembler la terre ? A l'origine des catastrophes sismiques*. Les Ulis, France: EDP Sciences.
- Boltanski, Luc, and Laurent Thévenot. 2006. *On Justification: Economies of Worth*. Princeton University Press.
- Cambrosio, Alberto, Peter Keating, Thomas Schlich, and George Weisz. 2006. 'Regulatory Objectivity and the Generation and Management of Evidence in Medicine'. *Social Science & Medicine* (1982) 63 (1): 189–99. <https://doi.org/10.1016/j.socscimed.2005.12.007>.
- Cardon, Dominique. 2016. 'Deconstructing the Algorithm: Four Types Types of Digital Calculations'. In *Algorithmic Cultures*, by Jonathan Roberge, Routledge.
- Daston, Lorraine. 1995. 'The Moral Economy of Science'. *Osiris*, 2nd Series, 10 (January): 2–24.
- Deleuze, Gilles. 2006. *The Fold*. New edition edition. London ; New York: Continuum.
- Desrosières, Alain. 2002. *The Politics of Large Numbers: A History of Statistical Reasoning*. Harvard University Press.
- Diakopoulos, Nicholas. 2014. 'Algorithmic Accountability Reporting: On the Investigation of Black Boxes'. <https://academiccommons.columbia.edu/catalog/ac:2ngf1vhhn4>.
- Domingos, Pedro. 2017. *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. 01 ed. Penguin Books Ltd.
- Dourish, Paul. 2016. 'Algorithms and Their Others: Algorithmic Culture in Context'. *Big Data & Society* 3 (2): 2053951716665128. <https://doi.org/10.1177/2053951716665128>.
- Hacking, Ian. 1990. *The Taming of Chance*. 1 edition. Cambridge England ; New York: Cambridge University Press.
- . 1999. *The Social Construction of What?* Harvard University Press.
- Hofman, Jake, Amit Sharma, and Duncan Watts. 2017. 'Prediction and Explanation in Social Systems'. *Science* 355 (February). <https://www.microsoft.com/en-us/research/publication/prediction-explanation-social-systems/>.
- Hourdeaux, Jérôme. 2016. 'Police Prédictive: Deux Chercheurs Démontent l'algorithme'. *Mediapart*, 2016. <https://www.mediapart.fr/journal/international/130916/police-predictive-deux-chercheurs-demontent-l-algorithme>.
- Ingold, Tim. 2013. *Making: Anthropology, Archaeology, Art and Architecture*. Routledge.
- Latour, Bruno. 1987. *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press.
- . 2010. 'Prendre le pli des techniques'. *Réseaux* 163 (5): 11–31. <https://doi.org/10.3917/res.163.0011>.
- . 2012. *Enquête sur les modes d'existence: une anthropologie des modernes*. La Découverte.
- Latour, Bruno, and Couze Venn. 2002. 'Morality and Technology'. *Theory, Culture & Society* 19 (5–6): 247–60. <https://doi.org/10.1177/026327602761899246>.
- Latour, Bruno, Steve Woolgar, and Jonas Salk. 1992. *Laboratory Life – The Construction of Scientific Facts*. Reprint. Princeton, N.J: Princeton University Press.
- Mackenzie, Donald. 2004. *Mechanizing Proof – Computing, Risk and Trust*. New Ed. Cambridge, Mass.: MIT Press.
- Main, Ian. 1999. 'Is the reliable prediction of individual earthquakes a realistic scientific goal?' 397 (1). <http://www.nature.com/nature/debates/earthquake/index.html?foxtrotcallback=true>.
- Manning, Peter K. 2008. *The Technology of Policing: Crime Mapping, Information Technology, and the Rationality of Crime Control*. NYU Press.
- Marsan, David, and Olivier Lengliné. 2008. 'Extending Earthquakes' Reach Through Cascading'. *Science* 319 (5866): 1076–79. <https://doi.org/10.1126/science.1148783>.

- Mittelstadt, Brent Daniel, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. 2016. 'The Ethics of Algorithms: Mapping the Debate'. *Big Data & Society* 3 (2): 2053951716679679. <https://doi.org/10.1177/2053951716679679>.
- Mohler, G. O., M. B. Short, P. J. Brantingham, F. P. Schoenberg, and G. E. Tita. 2011. 'Self-Exciting Point Process Modeling of Crime'. *Journal of the American Statistical Association* 106 (493): 100–108. <https://doi.org/10.1198/jasa.2011.ap09546>.
- Mohler, G. O., M. B. Short, Sean Malinowski, Mark Johnson, G. E. Tita, Andrea L. Bertozzi, and P. J. Brantingham. 2015. 'Randomized Controlled Field Trials of Predictive Policing'. *Journal of the American Statistical Association* 110 (512): 1399–1411. <https://doi.org/10.1080/01621459.2015.1077710>.
- Ogien, Ruwen, and Christine Tappolet. 2009. *Les concepts de l'éthique: Faut-il être conséquentialiste?* Paris: Editions Hermann.
- Olazaran, Mikel. 1996. 'A Sociological Study of the Official History of the Perceptrons Controversy'. *Social Studies of Science* 26 (3): 611–59. <https://doi.org/10.1177/030631296026003005>.
- O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.
- Pease, Ken, and Andromachi Tseloni. 2014. *Using Modeling to Predict and Prevent Victimization*. Springer Science & Business Media.
- Perry, Walt L. 2013. *Predictive Policing: The Role of Crime Forecasting in Law Enforcement Operations*. Rand Corporation.
- Porter, Theodore M. 1996. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton University Press.
- Robbins, Joel. 2010. 'On the Pleasures and Dangers of Culpability'. *Critique of Anthropology* 30 (1): 122–28. <https://doi.org/10.1177/0308275X09360136>.
- Rouvroy, Antoinette, and Thomas Berns. 2013. 'Gouvernementalité algorithmique et perspectives d'émancipation'. *Réseaux* 177 (1): 163–96. <https://doi.org/10.3917/res.177.0163>.
- Sandvig, Christian, Kevin Hamilton, Karrie Karahalios, and Cedric Langbort. 2014. 'Auditing Algorithms: Research Methods for Detecting Discrimination on Internet Platforms'. *Data and Discrimination: Converting Critical Concerns into Productive Inquiry*, May. <https://doi.org/10.1177/0308275X09360136>.
- Shmueli, Galit. 2010. 'To Explain or to Predict?' *Statistical Science* 25 (3): 289–310. <https://doi.org/10.1214/10-STS330>.
- Simondon, Gilbert. 2012. *Du mode d'existence des objets techniques*. Paris: Editions Aubier.
- . 2016. *On the Mode of Existence of Technical Objects*. Univocal Publishing LLC.
- Teil, Genevieve, and B. Latour. 2017. 'The Hume Machine. Can Association Networks Do More than Formal Rules?' *Stanford Humanities Review* 4, 4 (2, 2). <http://prodinra.inra.fr/?locale=fr#!ConsultNotice:153889>.
- Thomas van Stiphout, Jiancang Zhuang, and David Marsan. 2012. 'Seismicity Declustering'. <http://www.corssa.org>.
- Walker, Samuel. 2004. 'Science and Politics in Police Research: Reflections on Their Tangled Relationship'. *The ANNALS of the American Academy of Political and Social Science* 593 (1): 137–55. <https://doi.org/10.1177/0002716203262699>.
- Weisburd, David, Wim Bernasco, and Gerben J. N. Bruinsma. 2009. *Putting Crime in Its Place: Units of Analysis in Geographic Criminology*. Springer.
- Weisburd, David, and Anthony A. Braga, eds. 2006. *Police Innovation: Contrasting Perspectives*. Cambridge, UK ; New York: Cambridge University Press.
- Zarsky, Tal. 2016. 'The Trouble with Algorithmic Decisions'. *Science, Technology, & Human Values* 41 (1): 118–32. <https://doi.org/10.1177/0162243915605575>.
- Ziewitz, Malte. 2016. 'Governing Algorithms: Myth, Mess, and Methods'. *Science, Technology, & Human Values* 41 (1): 3–16. <https://doi.org/10.1177/0162243915608948>.

¹ PredPol is a small Santa Cruz start-up developed along the same path as many other California businesses. In 2010, two entrepreneurs—Caleb Baskin and Ryan Coonerty (also Third District Supervisor for Santa Cruz County)—approached two California researchers, George Mohler (associate professor of applied mathematics) and Jeffrey Brantingham (an archaeologist specializing in the Upper Paleolithic in northern China, and the son of two well-known criminologists who did pioneering work on the geography of crime), with a view to converting the fruits of their research into a profitable business with a strong growth potential. Although the research that underpinned PredPol was publicly funded, the start-up was created with 1.3 million dollars invested in 2012 by a handful of business angels. Following a business process that proved itself in spectacular fashion within two years, owing in particular to the efforts of its lobbyists operating in the Democratic networks of California, the firm was launched in a second round of venture capital fundraising (2.4 million dollars raised in 2014) in order to extend its commercial activity. At the time of writing PredPol is a commercial web-based system deployed in a number of policing departments in the United States and the UK.

² David Marsan's professional web page: https://www.isterre.fr/annuaire/pages-web-du-personnel/david-marsan/?id_auteur=131

³ I exchanged with David Marsan several times between February 2015 and March 2017. The quotations are drawn from my fourth interview with him in Chambéry in April 2016.

⁴ I made Marsan's analysis public in an article published on the web site of the magazine *Vie des idées* (Benbouzid 2016), as subsequently reported by *Mediapart* (Hourdeaux 2016).

⁵ Echoing Deleuze, Latour encourages us to *unfold* the technical action: "I would like to define the regime proper to technology by the notion of fold, without giving it all the Leibnizian connotations that Gilles Deleuze (1993) has elaborated so well. What is folded in technical action? Time, space and the type of actants" (Latour, 2002 p.248). In a methodological point of view, it means observing and following empirically what the PredPol's algorithm is concretely made of, in order to make its components visible.

⁶ PredPol's launch strategy was largely based on this slogan. On the basis of this marketing slogan, the public relations arm of PredPol sought to win over police leaders by convincing them that the PredPol system represents an improvement on "hotspot policing," one of the ways that proactive policing has been labeled since the 1990s. On their website (www.predpol.com), trials of the program are systematically associated to a decrease in crime of around 20%, a larger decrease than in sectors where the program is not used.

⁷ Marsan and his collaborator Lengliné published a widely noted article in statistical seismology in *Science* in 2008 (ibid.) This is the article that the PredPol mathematicians cite in their own article (ibid 2011). The statistical method that Marsan developed was integrally transposed, aside from the elements of translation that were essential to adapt the algorithm to the constraint of operationalization.

⁸ To represent the spread of victimization in statistical terms and to identify the more or less repetitive spatio-temporal configurations on which prevention strategies could be built, researchers apply spatial analysis statistical tools from epidemiological research. The analogy of contagion comes in 1990s with the notion of "communication of the risk of victimization". For example, the mechanism of contagion corresponded fairly well to the results of qualitative surveys run on burglars. Burglars had told researchers that they regularly returned to burgle the same house when it was easy to burgle and they had not been able to take everything the first time around. Burglars moreover operate by neighbourhood. The notion of "infectious burglaries" has been used to explain why victimization spreads in time and space. (Pease and Tseloni ibid.)

⁹ For a history of the controversial origins of neural networks, see Olazaran (1996).

¹⁰ The opacity of algorithms has become a commonplace observation: not only is the source code of machines usually protected as a trade secret, but it also describes a process of artificial learning that is so complex and that involves so many variables that the results are difficult to interpret even for the specialists themselves. Whether intentional or not, opacity is understood as a central problem in current public debate in all countries where the problem of algorithms has reached the political agenda (Mittelstadt et al. 2016).

¹¹ Marsan's complete critical note can be consulted through the article published on online news web site *Mediapart*, "Police prédictive: deux chercheurs démontent l'algorithme" [Predictive policing: two researchers take apart the algorithm], 13 September 2016.

¹² As Daston emphasizes, whereas "accuracy concerns the fit of numbers or geometrical magnitudes to some part of the world and presupposes that a mathematical model can be anchored in measurement [...] precision concerns the clarity, distinctness, and intelligibility of concepts, and, by itself, stipulates nothing about whether and how those concepts match the world." (ibid).

¹³ Although these aftershocks threaten the safety of rescue workers searching through ruins for survivors, Marsan's research does not aim to improve the prediction of their occurrence.

¹⁴ Interview with Sean Malinowski, August 2013.

¹⁵ The scientists at the startup undertook a serious evaluation of efficacy of the algorithm's algorithms, in the tradition of quasi-experimental methods (Mohler et al. 2015), which may seem to suggest that PredPol's approach is not deontological. Measuring efficacy is indeed a way of judging police actions in terms of results. However, the expRegarding this concept of robustness ntation test of PredPol was punctual: The cities that use the platform do not systematically carry out a randomized trial experiment. The assesement appears in the continuity of PredPol's marketing plan, rather than in a logic of "regulatory objectivity" (Cambrosio et al. 2006) as in seismology with the program "Collaboratory for the Study of Earthquake Predictability". <http://www.cseptesting.org/>.

¹⁶ The assessment implemented by PredPol appears more in the continuity of the start-up's marketing plan than in a logic of "regulatory objectivity" (Cambrosio et al., 2006) as in seismology with the "Collaboratory for the Study of Earthquake Predictability (CSPEP)" programme. <http://www.cseptesting.org/>, accessed 25/10/2017.

¹⁷ Robbin's analysis sheds light on the profession of "technological evangelist" (or Chief Evangelist Officer), which emerged in the world of information technology.