

Statistical Modelling of Pre-Impact Velocities in Car Crashes

*Wouter Kager** *Ivan Kryven[†]* *Keith W. Myerscough[‡]*
Timo van Opstal[§] *Thomas Rot**

Abstract

The law wants to determine if any party involved in a car crash is guilty. The Dutch court invokes the expertise of the Netherlands Forensic Institute (NFI) to answer this question. We discuss the present method of the NFI to determine probabilities on pre-impact car velocities, given the evidence from the crash scene. A disadvantage of this method is that it requires a prior distribution on the velocities of the cars involved in the crash. We suggest a different approach, that of statistical significance testing, which can be carried out without a prior. We explain this method, and apply it to a toy model. Finally, a sensitivity analysis is performed on a simple two-car collision model.

1 Introduction

Car crashes are unfortunately still common occurrences on the busy roads of the Netherlands. While it is of course preferable to prevent crashes from happening, as long as they do occur, it is important to determine if any party involved can be considered *guilty*. To determine this, Dutch courts invoke the expertise of the Netherlands Forensic Institute (NFI). The most common question asked in these situations is whether either party was driving too fast.

The NFI must answer this question based on data acquired at the crash scene by the relevant authorities. This is the *evidence*. To determine the velocities of both vehicles before the crash based on this evidence, two problems must be faced. Firstly, the physics of a vehicle collision has to be modelled accurately. There exists software which does this, but at the moment only forwards in time (see section 1.2). Secondly, the evidence is usually incomplete and has limited accuracy, as is discussed in section 1.1. The NFI tackles these issues with a Monte Carlo approach, briefly described in section 2, using the software PC-Crash. In this way an attempt is made to estimate the conditional distribution of the pre-crash velocity of a vehicle, given the evidence.

We argue that estimating this distribution requires a prior distribution on the pre-crash velocity, cf. section 3. The choice of prior is open to debate. The NFI assumes

*VU University

[†]University of Amsterdam

[‡]Centrum voor Wiskunde en Informatica

[§]TU Eindhoven

a uniform prior, but explains to the court that their calculations can be repeated with a different prior, if desired. We present an alternative approach, that of statistical significance testing, in section 4, which can be carried out without assuming a prior. This approach associates p -values to hypotheses, but the interpretation of these values is not a trivial matter (see section 4.2), and is also dependent on prior information. On the positive side, both the NFI and Dutch courts of law have experience in using p -values. Section 5 studies an example of the procedure applied to a toy problem.

Finally, a separate sensitivity analysis is presented in section 6. This analysis is performed on a simple two-car collision model with energy dissipation. Despite the simplicity of this model, it gives valuable insight into the sensitivity of the pre-crash velocity to the post-crash data.

1.1 Evidence Found at the Crash Scene

Whenever a car crash occurs, local authorities are called in. If the crash did not occur in a place extremely important for traffic flow (on a highway, for example), the commanding officer collects data, e.g. takes pictures of the cars, measures their positions and marks any tire tracks. It is important to note that the NFI is never present on a crash scene. The amount and accuracy of data is dependent on the diligence and experience of the commanding officer.

This means that collected data for each crash is different. Not only human behaviour is a factor in this. Natural surroundings also play a major role. For instance, if the crash occurs next to a lake, one of the cars may end up in the lake giving a huge amount of uncertainty in the end position of the car. All in all, we remark that the evidence depends on the situation and can have large uncertainties in the different parameters. It is therefore essential to have a sound and lucid method for dealing with these uncertainties. The method of statistical significance testing presented in section 4 satisfies these criteria.

1.2 The Software PC-Crash

PC-Crash is a deterministic simulation model for car crashes. The model consists of two parts: a collision model and a propagation model. The user sets up a scenario, placing cars and entering parameters. The program runs the propagation model until the cars make impact. An instantaneous collision is calculated, after which the propagation model is used to calculate the post-impact trajectories. In this paper we do not consider the correctness of this simulation model. We assume that it is perfect. Other studies [5] have shown that the model is reasonably correct for uncomplicated collisions. By uncomplicated we mean that we do not have factors which are not modelled in PC-Crash. For example, PC-Crash does not have the possibility of modelling cars that fall apart. It is also known that speed bumps are poorly modelled in the software. In those cases we cannot use PC-Crash.

The procedure proposed below is not dependent on the specific simulation program that is used. If a better program is developed we can incorporate this program into the procedure without any problem.

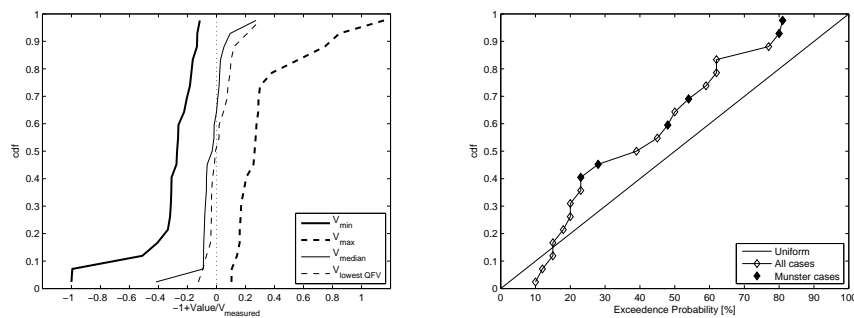


Figure 1: Left: the empirical cumulative distribution functions (cdf's) of the relative differences between the measured velocities and different ensemble velocities produced by MC-Crash. Right: the empirical cdf of the exceedance probabilities produced by MC-Crash.

2 Monte Carlo Simulation by the NFI

The approach developed by the NFI is a Monte Carlo procedure called MC-Crash, which is presented in [1]. An even more detailed set of instructions for researchers is given in [2]. [3] present some further information on the topic, but this does not represent the current methodology of the NFI. In summary, the NFI approach relies on random sampling from the input parameters (or “model parameters”) and measurement errors, and running PC-Crash for each parameter set. A quality function Q is used to determine how well a PC-Crash result matches the evidence. Private correspondence with Aart Spek, our contact at NFI, has taught us that the simulations are ended when the distribution of the 500 best parameter sets is visually close to that of the best 250, and that this usually requires 0.5–2 million samples. The PC-Crash output of the best 500 samples are stored.

In the NFI procedure, PC-Crash is treated as a black box function, which is assumed to model the physics involved correctly (cf. section 1.2). This has been verified by TNO at the request of the NFI [4, 5]. A validation of the Monte Carlo procedure of the NFI is presented in [4]. We now give a brief summary of this validation report.

The validation relies solely on comparing MC-Crash results to controlled test crashes. In these tests the pre-impact velocity was measured at the scene, making it possible to compare MC-Crash results with reality. As the output of MC-Crash is an ensemble of PC-Crash outputs, various statistical properties of this ensemble can be studied. For each test case, the velocity measured at the scene is compared with the minimum, maximum and median velocity of the MC-Crash ensemble, and also with the velocity of the sample which has the lowest Q . Finally, the measured velocity is compared with the so-called exceedance probability of the ensemble, which is the percentage of samples in the ensemble where the PC-Crash impact velocity exceeds the measured value.

The comparison between the MC-Crash results and the measured velocities can be visualised in two graphs (figures 4 and 5 in [4]), which we have reproduced in figure 1. In the validation report it is stated that “the exceedance probabilities . . . [are],

ideally, uniformly distributed between 0 and 100 percent”. Figure 1 gives us no reason to doubt that this is the case. As we can also see in this figure, for all test cases the measured impact velocities are between the minimum and maximum velocities produced by MC-Crash. Spek concludes that the median velocity and the velocity of the sample with the lowest Q “can be considered as estimators of the impact velocity”.

We remark that it is far from trivial how the choice of the quality function Q and the stopping criterion employed by the NFI influence the statistical properties of the final ensemble. Time restrictions prevent us from investigating these issues further. However, figure 1 does suggest that Monte Carlo simulation using PC-Crash can indeed give useful information about impact velocities from the crash scene evidence. In the following sections, we will explain what kind of conclusions we believe can (and cannot) be drawn from the evidence using simulation software such as PC-Crash without a prior distribution on the pre-crash velocities.

3 The Model

The software PC-Crash converts an initial state to post-impact trajectories, cf. section 1.2, and the evidence at the crash scene is a subset of this space. This is abstracted to the following diagram:

$$X \xrightarrow{\text{PC-Crash}} Y \xrightarrow{\text{filter}} Z . \quad (1)$$

The map PC-Crash is the simulation program. The space X is the space of input parameters and Y is the space of output parameters of PC-Crash. The latter space is extremely big, as it contains the complete post-crash trajectories of the cars involved. The space Z describes all measurable data at the crash scene. This data is incomplete, as was discussed in section 1.1. We model this by applying a filter on the output space Y . This filter retains only a small part of the information provided by the PC-Crash software, for instance only the end positions and orientations of the cars. It is important to note that Z will be chosen differently for different crashes.

We are not interested in what happens in the space Y , because we only have access to the data retained in Z from the crash scene. Therefore, we simplify the model once more. We denote the composition of PC-Crash and filter by $f = \text{filter} \circ \text{PC-Crash}$. What remains is the diagram

$$X \xrightarrow{f} Z . \quad (2)$$

The evidence e obtained from the crash is a point in Z , complemented with information about the “measuring error” in each of the measured components of e . This information can be used to construct a model for the measurement performed at the crash scene. If the actual input parameters for the crash would have been $x \in X$, then the outcome at the crash scene would have been the point $z = f(x) \in Z$. Under the assumption that x describes what actually happened, we then model the outcome of the measurement at the crash scene as the point $z + S$, where S is a random variable describing the uncertainty in the measurement. We assume that this measurement error is drawn from a distribution μ , which does not depend on z and is centred around 0. The conditional distribution of the measurement $z + S$, given that the outcome of the crash is z , is then obtained by shifting μ to the centre z . We denote it by μ_z . It is up

to the experts to decide which distribution describes the measurement error best. In a typical situation (think of measuring the end position of a car), it might be given by independent Gaussians in each coordinate.

We stress here that we do not have enough information to completely model the measurement at the crash scene: we only know its *conditional* distribution *given* that the crash had a particular outcome. As a consequence of this, we can make probabilistic statements about the evidence found at the crash scene, *under the assumption* that a particular pre-crash (input) scenario occurred. In the next section, we will discuss in more detail how to do this. However, we would like to turn things around and draw conclusions about the likelihood of different pre-crash scenarios given the evidence. It is important to realise that this will be impossible without *prior* information on the likelihood of different input scenarios (or a prior distribution on the space X). We will return to this issue in section 4.2.

4 Statistical Significance Testing

In a car crash involving two cars, denoted A and B , we might typically be interested in the following type of hypotheses:

1. A was speeding, but B was not;
2. B was speeding, but A was not;
3. A and B were both speeding;
4. Neither A nor B was speeding.

We will consider such hypotheses in the classical framework of statistical hypothesis testing. The idea is that we compute the probability of the measurement at the crash scene being at least as extreme as what we have actually observed, under the assumption that our hypothesis is correct. This probability is called the p -value. A very small p -value could give rise to the rejection of the hypothesis, since under the assumption that the hypothesis is correct, one would see the observed data only with a very small probability. Often, one takes as “null-hypothesis” the negation of the hypothesis one wants to show to hold true. If its p -value is sufficiently small, one rejects this negation, and this is evidence supporting the hypothesis one is interested in.

4.1 The Formalism

Suppose that we want to test a certain hypothesis $H_0 \subset X$, in the setting of the model described in section 3. Here one could think for instance of an H_0 containing all points in X where the initial speed of car A is smaller than some given speed v . The idea is now that we create a subset C of Z , called the *critical region*, such that finding the evidence in C can be considered unlikely under the hypothesis H_0 . It is of crucial importance that we define C independently of the actual outcome of the performed measurement, otherwise we would always be able to arrange things such that we could draw the conclusion we were after. So, to construct the critical region we have to use our model of the measurement errors, but we are not allowed to use the evidence e .

In order to define C we use an observable quantity, called a *test statistic*. For simplicity, we describe the procedure here for the case where the measurement errors are modelled assuming independent Gaussian errors in all coordinates, but the principle can be extended to cases where the measurement errors are modelled in a different way (although it might be more debatable what is the correct test statistic to use then). So suppose that $Z = \mathbb{R}^n$ and that μ is a normal distribution centred at 0, with standard deviation σ_i in the i -th coordinate, so that $\mu_{f(x)}$ becomes a normal distribution with the same standard deviations in all coordinates, centred at $f(x)$. Then, under the assumptions that the initial state was x , a proper choice for the test statistic would be

$$T_x(z) = \sum_{i=1}^n \frac{(z_i - f_i(x))^2}{\sigma_i^2}. \quad (3)$$

This choice is motivated by the fact that the level sets of the density of $\mu_{f(x)}$ correspond to the level sets of T_x .

Now we can define the critical region on level δ :

Definition 1. Fix $\delta > 0$. Define

$$C_x^\delta = \{z \in Z : T_x(z) \geq \delta\}. \quad (4)$$

The critical region on level δ for the hypothesis H_0 is the set

$$C^\delta = \bigcap_{x \in H_0} C_x^\delta = \{z \in Z : T(z) \geq \delta\}, \quad (5)$$

where $T(z) = \inf_{x \in H_0} T_x(z)$ is the test statistic under the hypothesis H_0 .

The interpretation of these critical regions is fairly straightforward. The test statistic T_x measures the “distance” of the measurement to a proposed outcome $f(x)$ of the crash. If the measurement is too far away from this proposed outcome, i.e. if T_x is too large, then we do not believe that the actual outcome of the crash could have been $f(x)$, hence that the actual input parameters could have been x . A complication arises when we want to make a statement about a composite hypothesis H_0 consisting of multiple points in the input space X . In these cases, to reject H_0 , we require that the measurement is far from $f(x)$, no matter which $x \in H_0$ we consider. Hence the intersection over $x \in H_0$ in (5). In terms of the test statistic T , the critical region C^δ is precisely the region where the test statistic exceeds the value δ , since $T(z)$ is the minimal value of $T_x(z)$ attained for any $x \in H_0$.

If we have evidence found at the crash scene, that is, a point $e \in Z$, then we can search for the smallest δ such that $e \in C^\delta$:

Definition 2. The critical δ^* for H_0 is defined by

$$\delta^* = \inf\{\delta > 0 : e \in C^\delta\} = T(e), \quad (6)$$

and the p -value of H_0 is the maximal conditional probability that the outcome of the measurement would lie in C^{δ^*} , under the condition that the initial state was any point $x \in H_0$. That is, formally,

$$p = \sup_{x \in H_0} \mu_{f(x)}(C^{\delta^*}). \quad (7)$$

In practice, p -values—being an upper bound—are only useful when they are small, i.e. when e is far from $f(H_0)$. The above procedure, however, can be followed for any null-hypothesis H_0 . It will assign a (useless) p -value of 1 to H_0 if $e \in f(H_0)$ (in this case $\delta^* = 0$, thus $C^{\delta^*} = Z$ from which it follows that $p = 1$).

4.2 Interpretation

In words, the p -value is the maximal “probability of exceeding the evidence” among the elements in H_0 . A small p -value suggests that H_0 is not very likely, because it means that the evidence found at the crash scene is far from what you would expect to have measured under any possible pre-crash situation contained in H_0 . More precisely, conditioned on any point $x \in H_0$ being the pre-crash state, the probability of the measurement being at least as far from $f(x)$ as the evidence found is less than the p -value.

Note that we have not phrased our conclusions in terms of the probabilities of the various hypotheses. To do so would require prior knowledge which will—we believe—typically not be available or disputable. This has consequences when one wants to compare different hypotheses. Suppose, for example, that the hypothesis H_0 has a much smaller p -value than the hypothesis H_1 . This gives us more reason to reject H_0 than to reject H_1 . However, if the prior probability of H_0 is much higher than the prior probability of H_1 , it might well be the case that H_0 and H_1 are actually roughly equally likely given the evidence.

Likewise, we might consider the hypotheses H_v that a given car A was driving at a speed of at most v just before the crash, and compare the p -values of these hypotheses as a function of v . If these p -values stay very small for low v , and increase strongly around a specific speed v_0 , this can be seen as evidence supporting the presumption that the car was driving at a speed close to v_0 . However, without a prior on the pre-crash speeds, we cannot turn this into a probabilistic statement about the distribution of the car’s speed before the crash.

Indeed, the framework for comparing two hypotheses is different, and makes use of the so called likelihood-ratio: one wants to compare the likelihoods, or probabilities, of H_0 and H_1 given the evidence $e \in Z$. Assuming for a moment that both H_0 and H_1 consist of only one point (x and y , respectively), we can use the continuous version of Bayes’ rule and obtain

$$\frac{P(H_0|e)}{P(H_1|e)} = \frac{g_{f(x)}(e)}{g_{f(y)}(e)} \cdot \frac{P(H_0)}{P(H_1)},$$

where $g_{f(x)}$ is the density of the distribution $\mu_{f(x)}$, $g_{f(x)}(e)/g_{f(y)}(e)$ is the “likelihood ratio”¹, and $P(H_0)/P(H_1)$ the prior odds for H_0 against H_1 . These prior odds are not part of the model. To determine them, we need for instance to know something about the overall (prior) probability that a car is speeding. Given that we need the prior odds, it is clear that this does yield probabilities, but the interpretation is questionable unless one can remove doubts about the prior assumptions.

¹When H_0 and H_1 are composite, that is, consist of more than one point, the likelihood ratio is not immediately well-defined, but bounds can be obtained by considering maximal or minimal values of $g_{f(x)}(e)$. Especially when H_0 or H_1 is extended, so that $f(H_0)$ or $f(H_1)$ contain both points close to and far from e , this yields very impractical and useless numbers.

5 An Example of Significance Testing

To illustrate the hypothesis testing outlined in section 4, we consider a simple model problem. A car with mass m travelling with an absolute velocity $v > 0$ crashes into a wall and loses a certain amount of the energy E_d due to vehicle deformation. After the collision it has an absolute velocity $z \geq 0$, thus $Z = \mathbb{R}_{\geq 0}$. Knowing v and E_d , we are able to compute the post-collision speed z by a model $z = f(v, E_d)$. The function f is easily computed using conservation of energy (momentum is not conserved due to the external forces of the wall), yielding

$$f(v, E_d) = \sqrt{v^2 - 2E_d/m}. \quad (8)$$

Note that in this simple model we see the deformation energy as a parameter, i.e. an element of the space X . In reality this is not the case: material parameters along with the impact speed determine the dissipated energy. We should note that the space X is not \mathbb{R}^2 , but $X = \{(v, E_d) \in \mathbb{R}^2 : v^2 - 2E_d/m \geq 0\}$. Not all deformation energies and velocities can occur together, since the deformation energy must be less than or equal to the kinetic energy corresponding to the velocity at which the car was driving, i.e. the car cannot lose more energy in the collision than it had in the first place.

We now assume that the post-crash velocity e is measured with an error, where the error has the normal distribution with standard deviation σ and mean 0. We are interested in computing the p -value of the null-hypothesis $H_0 = \{v < \tilde{v}, E_d^a < E_d < E_d^b\}$. This hypothesis assumes that before the crash the car was travelling with a velocity less than \tilde{v} and that the deformation energy was between E_d^a and E_d^b . The hypothesis is non-empty if $\tilde{v}^2 - 2E_d^a/m \geq 0$, which we will assume in what follows.

We are able to compute the p -value using formula (7). Due to the simplicity of the model we can do this analytically. In general this is not possible, and numerical integrators and minimisers should be used instead. We start by computing δ^* :

$$\delta^* = \inf \{ \delta > 0 : e \in C^\delta \}. \quad (9)$$

This is equivalent to computing $T(e)$:

$$\begin{aligned} \delta^* = T(e) &= \inf_{(v, E_d) \in H_0} \frac{(e - f(v, E_d))^2}{\sigma^2} \\ &= \begin{cases} \frac{(e - f(\tilde{v}, E_d^a))^2}{\sigma^2} & \text{if } e > f(\tilde{v}, E_d^a) \\ 0 & \text{if } e \leq f(\tilde{v}, E_d^a) \end{cases}. \end{aligned} \quad (10)$$

The last computation used the fact that $f(H_0) = [0, f(\tilde{v}, E_d^a)]$, because f is increasing in v , and decreasing in E_d . If $e > f(\tilde{v}, E_d^a)$, the infimum is attained at $f(\tilde{v}, E_d^a)$. If $e \leq f(\tilde{v}, E_d^a)$, then there is a point $(v, E_d) \in H_0$ such that $f(v, E_d) = e$, thus $\delta^* = 0$. As discussed in section 4, it follows that $p = 1$ if $e \leq f(\tilde{v}, E_d^a)$. It remains to compute the p -value when $e > f(\tilde{v}, E_d^a)$. In this case, from a simple computation it follows that $C^{\delta^*} = \{z \in Z : z > e\}$. Thus we compute using equation (7),

$$\begin{aligned} p &= \sup_{(v, E_d) \in H_0} \mu_{f(v, E_d)}(C^{\delta^*}) \\ &= \sup_{(v, E_d) \in H_0} \frac{1}{\sqrt{2\pi}\sigma^2} \int_e^\infty \exp\left(-\frac{(z - f(v, E_d))^2}{2\sigma^2}\right) dz. \end{aligned} \quad (11)$$

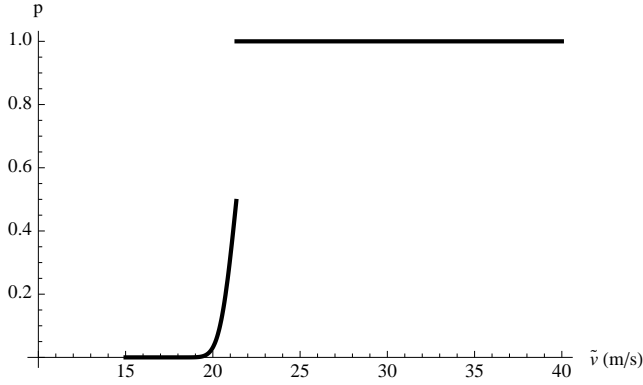


Figure 2: A plot of the p -values of the hypotheses $H_0 = \{v < \tilde{v}, E_d > E_d^a = 10^5 \text{ J}\}$. The measured post-crash velocity is $e = 16 \text{ m/s}$. The standard deviation is $\sigma = 1 \text{ m/s}$. The mass of the car is $m = 1000 \text{ kg}$. In other figures we show the response of the p -values to changes in the parameters σ , E_d^a , and the measured post-crash velocity e . The location of the sharp increase in the p -value suggests that the pre-crash speed was around this velocity.

The integral is easily expressed in terms of the error function:

$$\begin{aligned} p &= \sup_{(v, E_d) \in H_0} \frac{1}{2} \left(1 - \text{Erf} \left(\frac{e - f(v, E_d)}{\sqrt{2}\sigma} \right) \right) \\ &= \frac{1}{2} \left(1 - \text{Erf} \left(\frac{e - f(\tilde{v}, E_d^a)}{\sqrt{2}\sigma} \right) \right). \end{aligned} \quad (12)$$

In the last step we used that the function under the supremum is increasing in v and decreasing in E_d . We plug in the maximal v and minimal E_d in H_0 to obtain the result. By combining the previous computations, we conclude that

$$p = \begin{cases} \frac{1}{2} \left(1 - \text{Erf} \left(\frac{e - f(\tilde{v}, E_d^a)}{\sqrt{2}\sigma} \right) \right) & \text{if } e > f(\tilde{v}, E_d^a) \\ 1 & \text{if } e \leq f(\tilde{v}, E_d^a) \end{cases}. \quad (13)$$

We see that the p -value is insensitive to the upper bound on the deformation energy. This is due to the form of our null hypothesis. If we would compute the p -value for the null-hypothesis $H_0 = \{v > \tilde{v}, E_d^a < E_d < E_d^b\}$ we would find that the p -value is insensitive to the lower bound of the deformation energy.

In a realistic scenario it is impossible to compute the p -values analytically. We study some plots, obtained from the analytical result, that can be obtained in real scenarios with numerical computations using PC-Crash. In all plots we took the mass m of the car to be 1000 kg. We plotted the p -values of the hypothesis $H_0 = \{v < \tilde{v}, E_d > E_d^a\}$ against \tilde{v} in figure 2. The lower bound on the dissipated energy is $E_d^a = 10^5 \text{ J}$ and the standard deviation is $\sigma = 1 \text{ m/s}$. We see a sharp increase in the p -value just above $\tilde{v} = 20 \text{ m/s}$. At $\tilde{v} = 20 \text{ m/s}$ the p -value equals 3.2×10^{-2} , and at $\tilde{v} = 18 \text{ m/s}$ the p -value is 5.7×10^{-7} . As discussed in section 4.2, we therefore have a much larger confidence in rejecting the hypothesis that the car was driving at a velocity of at most 18 m/s than in rejecting the hypothesis that the car was driving at

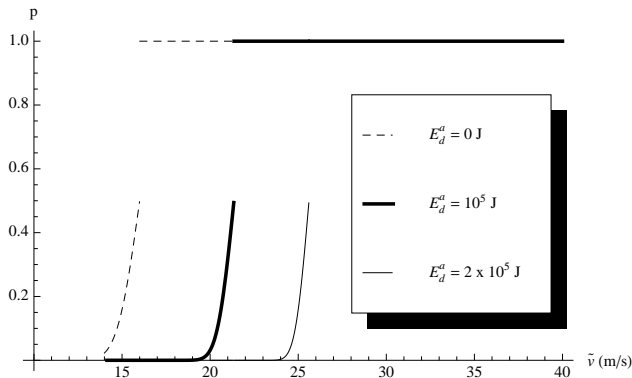


Figure 3: The plots depict the response of the p -values to varying dissipation energy. The p -values are plotted for different hypotheses $H_0 = \{v < \tilde{v}, E_d > E_d^a\}$ with varying E_d^a . The values of the other parameters are the same as in figure 2. We see that if we assume that more energy is dissipated, the increase in the p -values shifts (non-linearly cf. eq. (13)) to a higher pre-crash velocity.

a velocity of at most 20 m/s. However, without prior information we cannot state that, given the measured post-crash velocity, the probability of the car driving faster than 18 m/s is much larger than the probability of the car driving faster than 20 m/s.

Figure 3 shows the response of the p -values to changes in the value of E_d^a in the null-hypothesis. As expected, if we assume that more energy is dissipated, we are led to believe that the pre-crash velocity was higher. Figure 4 shows that an increasing uncertainty, i.e. a larger standard deviation, makes it harder to draw conclusions about pre-crash velocities. Finally, figure 5 investigates changes in the measured post-crash velocity. The behaviour is as expected: a higher post-crash velocity suggests that the car was driving faster before the crash. The discontinuities in all these plots occur at the smallest \tilde{v} such that $e \in f(H_0)$. From this point on the p -value equals 1, and becomes a trivial upper bound, hence is useless.

There is a small unmentioned subtlety in the previous computation. The post-crash velocity is a positive quantity and the probability distributions do have a finite probability mass for negative velocities. The assumption that the measurement error is normally distributed fails for small values of the velocity (theoretically, and practically as well). We have chosen to ignore this phenomenon, because the problematic values of the velocity (around $v = 0$ m/s) are many standard deviations away from the evidence.

6 Sensitivity-based Methods

This section contains a different approach to the problem. We consider a simple deterministic model for a collision of two cars. In this case we are able to invert the system of equations, i.e. we are able to compute the pre-crash velocities from the post-crash data. We are interested in understanding how small perturbations in the post-crash data affect the predicted pre-crash velocities. A powerful tool to analyse this are sen-

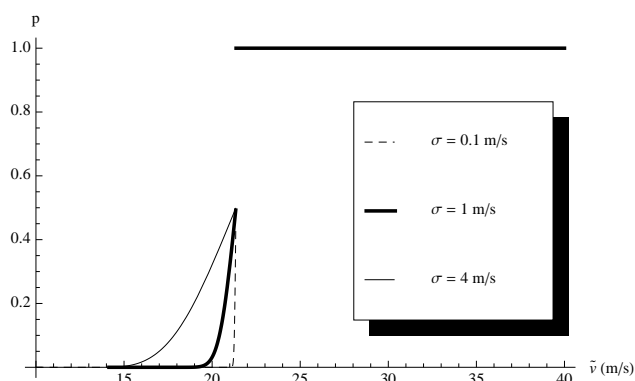


Figure 4: The plots depict the response of the p -values to a varying standard deviation σ . The p -values are plotted for different hypotheses $H_0 = \{v < \tilde{v}, E_d > E_d^a = 10^5 \text{ J}\}$, with varying standard deviation σ . Other parameters are the same as in figure 2. The increase in the p -values becomes less sharp with an increasing standard deviation, which makes it harder to draw strong conclusions.

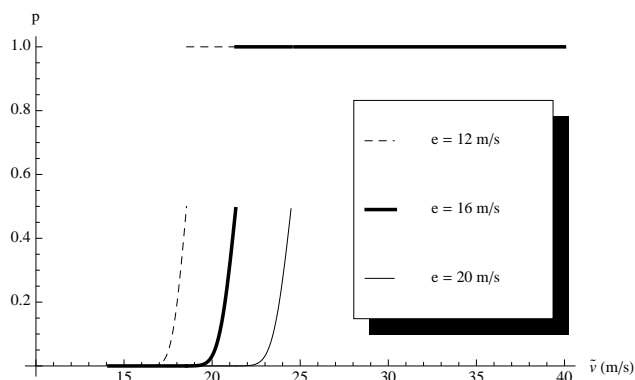


Figure 5: The plots depict the response of the p -values to a different measured post-crash speed e . The p -values are plotted for different hypotheses $H_0 = \{v < \tilde{v}, E_d > E_d^a = 10^5 \text{ J}\}$, with varying e . Other parameters are the same as in figure 2. If we measure a larger post-crash velocity, the increase in the p -values shifts to a higher pre-crash velocity, as expected.

sitivities. We will compute the sensitivities in this simple model and study them.

It should be noted that the sensitivities are ill-defined for collisions simulated using PC-Crash because the program is not invertible, i.e. we cannot uniquely reconstruct the pre-crash velocities from post-crash data. However, we do believe the discussion in this section gives valuable insight into the relative importance of measuring errors for the modelling of pre-crash velocities in real scenarios.

6.1 Deterministic Setting

We consider a simple collision model with two cars crashing into each other. We assume that the momentum exchange tangent to the collision is zero, so that the problem reduces to a one-dimensional collision. The cars obey the laws of momentum and energy conservation:

$$\eta u + \nu v = \eta u_+ + \nu v_+, \quad (14)$$

$$\frac{1}{2}\eta u^2 + \frac{1}{2}\nu v^2 = \frac{1}{2}\eta u_+^2 + \frac{1}{2}\nu v_+^2 + E_d. \quad (15)$$

In these equations (η, ν) , (u, v) and E_d denote vehicle masses, pre-impact velocities and impact dissipation energy, whilst (u_+, v_+) denotes the post-impact velocities.

This is a quadratic equation, and it is not difficult to compute the pre-crash velocities u, v , if we know the post-crash velocities u_+, v_+ , and the dissipated energy E_d :

$$u = (\nu v_+ + \eta u_+ \pm \nu \sqrt{D}) / (\nu + \eta) \quad (16a)$$

$$v = (\nu v_+ + \eta u_+ \mp \eta \sqrt{D}) / (\nu + \eta). \quad (16b)$$

The discriminant D is given by

$$D = (v_+ - u_+)^2 + 2 \frac{\eta + \nu}{\eta \nu} E_d. \quad (17)$$

The signs \pm depend on the data u_+ and v_+ . In general a one-dimensional collision model such as this one has two solutions. In one solution the colliding objects bounce off each other, and in the other solution they move through each other. Both collisions are governed by the same equations. An example of the latter type of solution is a bullet shooting through a tin can.

In our situation the cars do not move through each other. This means that if $v_+ > u_+$, we must have $u > v$, so we need the $+$ sign in equation (16a), and the $-$ sign in equation (16b). In the case $v_+ < u_+$ we must use the opposite signs. If $u_+ = v_+$, we cannot be sure which situation has occurred. This will cause a discontinuity at $u_+ = v_+ \neq 0$.

6.2 Sensitivities

For the remainder we analyse u . The situation for v is analogous. Small variations in the parameters u_+, v_+, E_d, η , and ν give rise to variations in u . This can be understood from the differential change in u :

$$du = \frac{\partial u}{\partial u_+} du_+ + \frac{\partial u}{\partial v_+} dv_+ + \frac{\partial u}{\partial E_d} dE_d + \frac{\partial u}{\partial \eta} d\eta + \frac{\partial u}{\partial \nu} d\nu. \quad (18)$$

We are interested in how these terms relatively influence du . We rescale all the terms and write

$$du = u_+ \frac{\partial u}{\partial u_+} \frac{du_+}{u_+} + v_+ \frac{\partial u}{\partial v_+} \frac{dv_+}{v_+} + E_d \frac{\partial u}{\partial E_d} \frac{dE_d}{E_d} + \eta \frac{\partial u}{\partial \eta} \frac{d\eta}{\eta} + \nu \frac{\partial u}{\partial \nu} \frac{d\nu}{\nu}. \quad (19)$$

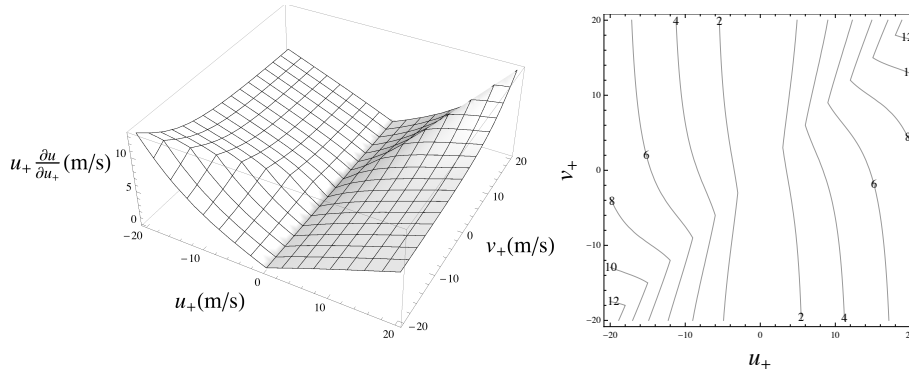


Figure 6: The sensitivity to u_+ is plotted, as a function of u_+ and v_+ . The other parameters were kept fixed: $\eta = 2000$ kg, $\nu = 1000$ kg, and $E_d = 5 \times 10^4$ J.

The equation expresses how u is modified by small relative changes in each of the parameters. Thus the absolute value of for instance $u_+ \frac{\partial u}{\partial u_+}$ determines the sensitivity of u to relative changes in u_+ . We call $u_+ \frac{\partial u}{\partial u_+}$ the sensitivity of u to u_+ . The sensitivities of u to u_+ , v_+ , and E_d are given by

$$\begin{aligned} u_+ \frac{\partial u}{\partial u_+} &= \left(\frac{\eta}{\eta + \nu} \pm \frac{\nu(u_+ - v_+)}{\sqrt{D}(\eta + \nu)} \right) u_+ \\ v_+ \frac{\partial u}{\partial v_+} &= \left(\frac{\eta}{\eta + \nu} \pm \frac{\nu(v_+ - u_+)}{\sqrt{D}(\eta + \nu)} \right) v_+ \\ E_d \frac{\partial u}{\partial E_d} &= \left(\pm \frac{1}{\eta \sqrt{D}} \right) E_d. \end{aligned}$$

In principle we could also have computed the sensitivities to η and ν , but we will focus on the sensitivities to u_+ and E_d . The sensitivity to u_+ is plotted in figure 6 as a function of u_+ and v_+ .

It is not difficult to obtain bounds for the sensitivities to the post-crash speed and dissipation energy from our explicit formulas. If we work under the assumption that $u_+^2 \geq v_+^2$, for instance, we can reason as follows. First we express the dissipation energy in terms of the post-crash kinetic energy of the first car by introducing a positive parameter α such that

$$E_d = \frac{4\alpha^2\nu}{\eta + \nu} \cdot \frac{1}{2}\eta u_+^2.$$

Since $(v_+ - u_+)^2 \leq 4u_+^2$, we can now bound the discriminant D by

$$4\alpha^2 u_+^2 \leq D \leq 4(1 + \alpha^2)u_+^2,$$

from which we obtain

$$\frac{\alpha^2}{\sqrt{1 + \alpha^2}} \frac{\nu}{\eta + \nu} |u_+| \leq \left| E_d \frac{\partial u}{\partial E_d} \right| \leq \alpha \frac{\nu}{\eta + \nu} |u_+|.$$




Figure 7: The ratio of sensitivities to u_+ and E_d is plotted, as a function of u_+ and v_+ , for different values of E_d . The top leaf is plotted for $E_d = 1 \times 10^4$ J. The middle leaf has a dissipation energy $E_d = 5 \times 10^4$ J, and the bottom leaf has $E_d = 2.5 \times 10^5$ J. For small E_d the sensitivity to u_+ is important, and for high E_d the sensitivity to E_d becomes more significant.

On the other hand, $(v_+ - u_+)^2 \leq 4u_+^2$ also implies that

$$D \geq (1 + \alpha^2)(v_+ - u_+)^2,$$

from which we conclude that

$$\frac{\eta - \nu/\sqrt{1 + \alpha^2}}{\eta + \nu} |u_+| \leq \left| u_+ \frac{\partial u}{\partial u_+} \right| \leq \frac{\eta}{\eta + \nu} |u_+|.$$

These bounds show that if the masses of the two vehicles are of the same order, the sensitivity to the dissipated energy will be larger than the sensitivity to the post-crash speed (of the first vehicle) when α is large. When α is small, on the other hand, the sensitivity to the dissipated energy stays small as well, and the sensitivity to the post-crash speed will be larger (provided the vehicle masses are not the same). It is to be noted here that large or small α correspond respectively to the dissipated energy being much larger or much smaller than the post-crash kinetic energy of the first vehicle. Thus, we can conclude that which sensitivity dominates depends on how the dissipated energy compares to the post-crash kinetic energy of the vehicles. This is visualised in figure 7.

6.3 Stochastic Analysis

In the simple toy model studied above, we could explicitly solve for the pre-crash speed u , given the output parameters (u_+, v_+, E_d) . If we are in such a situation, we may map the measured outcome of the crash directly to a pre-crash speed u , and we may wonder how good an estimate of the pre-crash speed this will give us. It is this question which we will address here.

To generalise the problem, we assume in this discussion that the output space Z is d -dimensional, and that every point $z = (z_1, \dots, z_d) \in Z$ maps uniquely to a pre-crash speed $u(z)$. The toy model discussed above fits in this picture if we take $d = 3$ and identify the coordinates (z_1, z_2, z_3) with the respective output parameters (u_+, v_+, E_d) . As in section 3, if we now assume that the actual outcome of the crash

is given by the point $z \in Z$, we can model a measurement of the outcome of the crash as the point $z + S$, where S is a random variable describing the measurement error. As before, we assume that this error is drawn from a distribution which is centred around 0.

To simplify the discussion, we furthermore assume the following:

1. The coordinates of S are independent. In this case the probability density function g of S is simply the product of the densities of its components. This is a simplification for presentation purposes, but not an essential condition for the method presented here.
2. The density of each coordinate of S is symmetric around 0, and the i -th coordinate has standard deviation σ_i .
3. The pre-crash speed depends sufficiently smoothly on the data around the given point z . For instance, in the example discussed above we require the relation (16a) to be sufficiently smooth in a neighbourhood of (u_+, v_+, E_d) we are interested in.

The 3rd-order Taylor series expansion of $u(z + s)$ around z is given by

$$u(z + s) = u(z) + \sum_{i=1}^d \frac{\partial u(z)}{\partial z_i} s_i + \frac{1}{2} \sum_{i=1}^d \sum_{j=1}^d \frac{\partial^2 u(z)}{\partial z_i \partial z_j} s_i s_j + O(|s|^3). \quad (20)$$

In the setup under discussion, each output point is associated in a deterministic way to a corresponding pre-crash speed. Thus we can map our measurement $z + S$ to the pre-crash speed $U = u(z + S)$. This new random variable describes which conclusion we would draw about the pre-crash speed from a measurement at the crash scene using our deterministic model, *under the condition* that z was the actual outcome of the crash (hence $u(z)$ was the actual pre-crash speed). If we know the map explicitly, we can now study properties of this random variable U , and determine how good an estimate of the actual pre-crash speed it will be. In particular, we can analyse how this estimate depends on the measurement errors in the different coordinates.

From the definitions of stochastic moments, for instance, the mean m_U and standard deviation σ_U of the random variable U are given by

$$m_U := \int_Z u(z + s) g(s) ds \approx u(z) + \frac{1}{2} \sum_{i=1}^d \left(z_i^2 \frac{\partial^2 u(z)}{\partial z_i^2} \right) \left(\frac{\sigma_i}{z_i} \right)^2; \quad (21a)$$

$$\sigma_U^2 := \int_Z (u(z + s) - \mu_U)^2 g(s) ds \approx \sum_{i=1}^d \left(z_i \frac{\partial u(z)}{\partial z_i} \right)^2 \left(\frac{\sigma_i}{z_i} \right)^2. \quad (21b)$$

We have expressed these moments in terms of the relative errors σ_i/z_i in each coordinate, and we then see the sensitivities to the different coordinates pop up in these expressions. The errors on these approximations are 4th order in the relative errors, because under the assumptions made above, the odd moments of the measurement errors vanish and different coordinates are independent. We see that the mean shifts according to the second derivatives of u , and that the standard deviation changes according to the squares of the sensitivities.

7 Conclusion

In this paper we propose an alternative approach for determining pre-crash velocities from crash scene evidence using an available black-box crash model. This *statistical significance testing* yields p -values for a set of pre-formulated statistical hypotheses. These are valuable indicators frequently used in modern statistical analysis and forensic science. It should be noted that there are no limitations on the complexity of the collision model. Any (third party) software is acceptable.

An example using a simple crash model was presented to help understanding the nature of uncertainty in the reconstructed data and how this is reflected by the p -values. This procedure can be adapted to use any other crash model, such as PC-Crash.

A sensitivity analysis was performed on a simple two-body collision. It was observed that sensitivity to energy dissipation depends on the amount of dissipated energy. It is particularly important to measure the dissipated energy with a high relative accuracy in collisions in which a large proportion of the kinetic energy is dissipated.

The presented approach provides a statistical foundation for a forensic expert analysing a car crash, and for presenting this data in court.

Acknowledgements

We would like to thank Geert Geeven, Ronald Meester, Thanh Son Nguyen, Bob Planqué, and Aart Spek for many fruitful discussions on this topic.

References

- [1] Reconstructie bots- en aanrijdsnelheden mbv MC-Crash. Technical Report 272501, NFI.
- [2] K.M. Hagendoorn. *Handleiding MC-Crash*.
- [3] A. Moser, H. Steffan, A. Spek, and W. Makkinga. Application of the monte carlo methods for stability analysis within the accident reconstruction software pc-crash. *Society of Automotive Engineers, Inc.*, January 2003.
- [4] A.C.E. Spek. Validatierapport MC-Crash. Technical report, NFI, January 2003.
- [5] H.H. Spit. Evaluation of PC-Crash - A Momentum-Based Accident Reconstruction Program. Technical Report 00.OR.BV.2712.1/HHS, TNO, 2000.