Estimating the Spatial Distribution of Environmental and Public Risk due to Weapons Firing Ranges

Alun Pope and Xun-Guo Lin
Department of Statistics
University of Newcastle
NSW 2308 Australia
Email: {pope, linx}@maths.newcastle.edu.au

Abstract: New procedures for the management of risk---to the public, the environment and military personnel---have been developed for weapons firing ranges, to replace ones which lacked an adequate conceptual and scientific basis. In this paper, these new techniques will be described, and examples given of their use. The computational basis of the techniques is simulation, but naive simulation is prohibitively expensive, so we shall discuss how simulation can be modified to obtain results, and especially how the necessity of obtaining estimates of the accuracy of the results affects the computational techniques. An important feature of our techniques is that we are able to compute estimates of the spatial (ie 2D, 3D, and also time-varying) distribution of the risk. The results can be presented in many ways, the choice of method depending on the problem at hand, or the aspect that it is desired to model. One important aspect is the immediate safety of the public; another, but to date less explored, aspect is the assessment of the spatial distribution of damage to the environment through destruction of flora (eg by explosion), and the dispersal of chemical contaminants (eg lead) in the range region.

1 INTRODUCTION

When the rockets go up,

Who cares where they come down? (Lehrer)

1.1 The range safety problem

How safe are weapons ranges? In particular, how does the risk of injury or damage depend on the location of a vulnerable object?

Many different types of weapon are fired on many different types of range. On military ranges the weapons fired may be automatic, such as machine-guns, deliver munitions with explosive charges, such as guided missiles and artillery shells, or munitions which release multiple submunitions (bomblets). Moving targets and firing platforms may be involved, as in air-to-ground, or air-to-air practices. The firing range may be intrinsically three-dimensional; the maximum range of the weapons may vary from a few to many kilometres; the damage caused may be localised (small arms) or widespread (high explosive shells).

Decisions about how large firing ranges should be, and what shape, involve choices about the use of resources and an assessment of the relative importance of competing components of the public good. Satisfactory policy-making under these circumstances requires political decisions and the resolution of complex technical issues.

1.2 The lack of a scientific basis

I shot an arrow into the air,
It fell to earth, I know not where. (Longfellow)

In the field, danger areas are often defined by laying a 'template' over a map of the range. (For an example, see Section 3.4.) The size and shape of the template depend upon the weapon being fired, the ammunition and the type of practice. As observed in Amies (1986), in Australia the official manuals refer to no mathematical or statistical analysis upon which the procedures could be based. The basis of many danger area template dimensions appears to lie entirely in a consensus of informally collected range experience. For example, Australian practice on small arms ranges was based at least until 1967 on a sequence of British War Office publications. It is said (Australian Delegation to IRSAG, 1989) that the minimum allowed height for firing live ammunition over the heads of troops originates from the British War Office's requirement of 1908 that this exceed by 6 feet the height of a lancer mounted on horseback.

The reasons for concern There have been serious accidents involving the public. In the United Kingdom, for example, in 1983 a woman was fatally wounded while near but outside the boundary of Ash Range in Surrey (The Times, 1983) and in 1988 an eleven-year-old boy was hit by shrapnel from a 105mm shell accidentally fired into the village of Enford (The Times, 1988). In other countries too there have been accidents involving members of the public outside range boundaries, but these are less publicly documented. From the historical record it cannot simply be *assumed* that firing ranges are safe.

It is the duty of those responsible for the management and operation of firing ranges to provide scientifically-based estimates of the degree of risk to which the public is exposed. We now have available a reasonably complete theoretical model of many aspects of the problem, on the basis of which it is possible to make predictions under realistic assumptions. (See for example Pope, 1988, 1991.) The Australian Ordnance Council has adopted this theoretical model in principle. However, at present, insufficient experimental work has been done, so that the data required to use the model are lacking. It is hoped that future international collaboration will rapidly rectify this situation.

2 A MATHEMATICAL MODEL

In addition to position, the motion of a projectile is determined by the values of other parameters, such as momentum of the centre of mass and angular momentum. The *state* of a projectile at a particular point in its motion is defined to be the vector consisting of the values of these extra parameters. It is assumed that a set of these parameters has been chosen which is of minimal dimension but is sufficient together with position to characterise the motion. The *phase* of the projectile is then the vector obtained by combining its position and state. It is a function of time.

2.1 Probabilistic definitions of risk

It is natural to assume that we are interested only in events whose occurrence depends *only* on characteristics of the projectile motion. We therefore choose to model only events which are determined as a function of the phase and the time: the probabi;ity density function (pdf) of such an event can in principle be obtained from the joint pdf of the phase and time. Events of this type have pdfs which can be traced back to the pdfs of the initial conditions (which presumably can be experimentally determined). Ricochet (or any other discontinuity in the phase as a function of time) disrupts this tidy scheme, but if the initial conditions of post-ricochet flights can be connected through a probabilistic model with the final conditions before ricochet, it turns out that it is possible to connect the pdf of the final impact points with the pdf of the initial conditions even after ricochets. This approach is adequate to treat a wide variety of questions of practical interest.

It is not enough to consider the pdf of the projectile merely in space and time, *even if* we are not interested in the other components of the phase, and will integrate them out to obtain the appropriate marginal pdfs. Because the trajectory is not uniquely determined by the requirement that it pass through a particular point in space, it becomes extremely complicated—perhaps impossible—to avoid double counting when we integrate the pdf to obtain probabilities. On the other hand, the defining property of the phase, that it *characterises* the motion, eliminates all double-counting problems and applies in greater generality.

2.2 The scrapmetal problem

Previous work on defining levels of risk concentrated implicitly on the **problem of containment**: find a closed region C on the range surface which contains a large proportion of the projectile impacts. Clearly this can lead to a reasonable definition of a safe area for persons standing on the ground.

Although not directly related to the issue of immediate safety (since spent bullets are not generally a hazard), the containment problem is significant because defining an area outside which only a minute fraction of bullets are likely to be found might give a politically acceptable method of defining range boundaries.

The main disadvantages of using this method to determine range boundaries are that it is conservative and ignores the airspace above the range. Nonetheless, it is worth computing, because it gives predictions which can be checked against observation and it allows a comparison to be made with existing safety boundaries.

Further, the solution to the containment problem we illustrate below also provides most of the answer to one of the environmental problems involved in range use: where has all the material gone that was fired down the range? We call this the **scrapmetal problem** (these materials may include heavy metals for example) and note that it also provides the ability to calculate levels of risk inside the range, with the intention of allowing the exposure of personnel in training to a greater (but known) level of danger.

We show in Section 3 that our methods can solve the scrapmetal problem in practice.

2.3 Components of the model

We regard the motion as consisting of a finite number of *flights* between *break points*. During flight, the motion is assumed to be governed by differential equations of known form. These concepts will be made more precise below.

The model has the following components:

1. break points (eg impacts on the range surface, where ricochet may occur);
2. a launch model, specifying the joint probability distribution of the launch parameters (eg from where, in what direction and how quickly the projectile is fired);
3. flight models, in the form of differential equations;
4. transition models (eg regression of outputs on inputs) providing the link between phase-time at termination of one flight (that is, at a break point) and the pdf of the initial conditions of the next flight, together with the probability that there is a next flight and the pdf of the coefficients in the flight model for the next flight; and
5. an environmental model for atmospheric conditions, range topography, etc.

It was demonstrated in Pope (1991) that the general problem could be solved. For example, each of the following can be computed:

- pdf of point of last impact on range surface
- probability of hitting a stationary object
- probability of hitting moving objects

- probability of hitting a moving object from a moving firing platform
- probability of hitting an object with sufficient energy to cause damage.

The solution covers important special cases, including:

- multiple and moving firers and targets;
- all combinations of air, ground and sea practices;
- any risk that can be defined in terms of position and momentum of the projectile;
- general range topography;
- a range surface constructed of non-uniform material;
- variation in behaviour due to weather.

In addition the risk can be computed at any point in space and at any time. The principal projectile type that is excluded is guided weapons.

3 A NUMERICAL EXAMPLE

Careful consideration needs to be given to the numerical methods employed in practice. The approach we have adopted in the example described below combines simulation and smoothing. Our problem is to compute the probability density function of last impact position (i.e. the scrapmetal problem).

3.1 The assumptions

The range surface and the launch distribution. We model the firing of 0.5 inch bullets from a prone position at a small target 550m away on a range which is flat apart from a hill in the form of a spherical cap centred at (600,0), which represents a stop butt at 600m. (Coordinates of points on the range are given as (x, z), denoting x metres downrange and z metres across, with the positive z direction being to the *right* of the positive x-axis.) The hill has a radius of 15m where the sphere protrudes through the plane of the rest of the range; the radius of the sphere is 20m; this means the hill has a height of 6.77m.

A fixed firing point is taken 0.3m above the range surface at (0,0). The elevation and azimuth angles of firing and the muzzle speed are assumed to have independent normal distributions. The mean firing elevation is 0.01 radians, which without the hill would result in a first impact on the range surface near 1000m. The mean firing azimuth is 0, which is along the positive x-axis. The standard deviations of elevation and azimuth are 0.007 and 0.011 radians, respectively. These values are consistent with those given in UK Ordnance Board (1989) as the values achievable in practice (with a different weapon) according to expert opinion. The mean muzzle speed is about 900 ms^{-1}, and its standard deviation is 10 ms^{-1}. This standard deviation is approximately that observed in tables (for a different ammunition) given in UK Ordnance Board (1984).

The above combination of angles was chosen because some bullets hit the stop butt and some clear it; in addition some strike it a glancing blow, which makes ricochet more likely. The shape of the stop butt is unusual although the height is typical.

Ricochet models. Ricochet was simulated from a combination of regression models fitted to experimental data for two different types of ammunition. The first component of this model was derived from data relating to 7.62mm ammunition, appearing in UK Ordnance Board (1984). It is appropriate for impact on turf, and predicts output elevation and velocity from input elevation and velocity, given that a ricochet has occurred. This prediction is combined with a predicted azimuthal deflection obtained from a regression model fitted to experimental data for 0.5 inch ammunition supplied by the US Army at Picatinny Arsenal (ARDEC, 1990).

Although these regression models combine data for different types of ammunition, the models employed are dimensionless: we can hope they represent to some extent the geometry of ricochet, which perhaps does not vary much between these two ammunitions. No better models are presently available.

Whether ricochet occurred at a given impact was decided by a very crude model fitted to the very small amount of data available (ARDEC, 1990) relating probability of ricochet to input variables such as angle and speed of impact, and the nature of the surface.

Caveat The reader will have observed that, while the modelling is detailed, where data values or appropriate models of behaviour were not known, they were *assumed*. However, the fact that values have been invented means that *it would be entirely inappropriate to use the results below for any purpose other than to evaluate the approach being adopted.*

3.2 The equations of motion and their solution

The equations of motion for a point-projectile in the absence of wind may be written:

$$\frac{d\mathbf{v}}{dt} = -\rho K(v)v\mathbf{v} - \mathbf{g}$$

where $\mathbf{v}$ denotes the velocity of the projectile, v is the magnitude of $\mathbf{v}$, $\mathbf{g}$ is the acceleration due to gravity, ρ is the density of the air (which may vary with position) and the coefficient $K(v)$ is determined as described below.

The determination of K depends (in our model) on whether the equation is being applied to stable or unstable flight. In stable flight, the projectile has low drag because its orientation is stable with the long axis approximately aligned with the direction of flight; but we take K to depend on speed relative to the air. In unstable flight the projectile tumbles or has a high yaw angle, with the long axis precessing about the line of flight, either of which cause the drag to be much greater. For unstable flight, the values of K were chosen to be appropriate for modelling the flight of fragments of irregular shape. The numerical values were obtained from Bentley (1986a,b).

The differential equations were solved using the subroutine lsodar from the library ODEPACK available from netlib. This is a variable steplength program which allows testing for a terminating condition.

3.3 Approach to the numerical calculation

A naive Monte Carlo simulation cannot be expected to provide the sort of result we want because we are particularly interested in evaluating the probability density at points well away from the centre of the distribution. Monte Carlo simulation by its nature concentrates on the most probable regions, so to gain adequate accuracy at the extremes of the distribution by simulation, a very large number of replications would be required.

The overall strategy we have adopted is a more sophisticated version of the Monte Carlo approach. It can be divided into two stages: simulation and smoothing. The output of the first stage is a histogram of impact probabilities, which is smoothed in the second stage.

3.3.1 Simulation Step

Systematic sampling. The approach described here is a simple variance reduction technique: we simulate ricochet but take a grid of values from the distribution of launch elevations and azimuths. The effects of ricochet are introduced in a Monte Carlo simulation, but the sampling of the launch distribution is systematic. By sampling systematically, sufficiently many points are obtained in the regions of sparse impacts to allow reasonable estimates to be made of the density there; at the same time, it is possible to keep the number of simulated rounds fired within reasonable limits. The systematic sampling forces points to be sampled from the extremes of the distribution.

The 'bullets' used in the simulation are given weights proportional to the value of the joint probability density function of the launch parameters evaluated at their particular combinations of muzzle speed, elevation and azimuth. When the impact-point of a particular bullet is calculated, the coordinates and the weight are recorded. When this has been done sufficiently often, the result can be thought of as a list of random coordinates and weights which can be displayed as a three-dimensional plot of weight against final resting point on the range surface. Storage considerations and convenience lead us to collect the weights in an array of bins on the range surface, so that this plot is a histogram. This histogram must then be smoothed and normalised to provide an estimate of the pdf.

Randomised systematic sampling. This approach is similar to systematic sampling but is motivated by the need to compute error estimates for our results. It is a version of importance sampling, involving taking initial conditions selected randomly from a known sampling distribution defined on the space of all possible parameter values. An example is given in Lin and Pope (1995).

3.3.2 Smoothing In Figures 1 to 4, we have simply formed a grid over the range surface and added the probability weights of impacts in each grid cell. This results in histogram estimates that are too smooth in some parts and too rough in others. What is really required is an adaptive technique, which smooths by different amounts depending on the local behaviour of the estimated density. However we have not (yet) attempted this.

3.4 Results

In Figure 1 is shown a perspective plot of $10\log_{10}(1+f)$ against position on the range surface, where f is the probability density function. The estimate is given again in a contour plot in Figure 2. The maximum range of the weapon is approximately 7000m, so the entire range can be seen. No impacts were recorded outside the plot-region. The smoothing in Figures 1 and 2 was carried out as indicated above by aggregating probability weights within cells with edges 140m down and 100m across range, forming a histogram on a rectangular grid.

It will be noticed that the risk is very concentrated in this example along a narrow corridor between the firing point and the stop butt. The true corridor is most likely narrower than indicated, as this part of the pdf is surely oversmoothed by aggregating on our rather coarse grid.

Spread out beyond the stop butt in a horseshoe shape are impacts due primarily to ricochet. It will be observed that there is a slight tendency for bullets to be found to the right of the mean line of fire. This asymmetry is due to the fact that the data used to provide the ricochet models came from firing real bullets, which spin; this spin apparently causes them to be deflected to the right on ricochet. While it is true that the trajectory of a spinning bullet also tends to curve (also to the right in this case), this effect is not included in the flight equations solved in this program as only a simple point-mass projectile model was used. Introducing this curvature would thus tend to increase the asymmetry, but only as a second-order effect.

Comparison with the safety template The polygon overlaid on the contour plot in Figure 2 shows the safety template for the ammunition being used. The intention in using a safety template is that it should contain all the fired bullets. It is clear from this Figure that, if the model is accurate, this is not the case. It is of interest therefore to calculate the probability of hitting an object outside the safety template. To achieve this, a square witness-box centred at (3400,1400) with sides parallel to the axes was introduced into the model. This box was introduced solely to record impacts. The box was 2m high, had sides of length 400m, and did not permit ricochet. The best estimate of the probability per round of hitting this box was 9.3×10^{-6}, which gives an average probability per unit area for the box of 5.8×10^{-11} m^{-2}. If one assumes that a box one metre square and 2m high represents a person, this suggests that the probability per round of hitting a person standing near (3400,1400) is 5.8×10^{-11}. This seems appropriately low, even when it is remembered that the ammunition being modelled is fired from a machine gun, so that millions of rounds per year may be fired using this safety template. Because of the length of time it takes to perform one such calculation, we have not been able to estimate the standard error of our estimate of this probability. (The estimate given is based on the simulated firing of 2.2×10^{6} bullets, and took three days to compute on a SUN Sparc1+ workstation.)

Effect of the butt on risk near but outside the safety template It is of interest to compare the above plots with what is obtained if the range is flat. This is shown in Figures 3 and 4. In comparing Figures 1 and 3, it should be noted that the horseshoe shaped hill of ricochets is approximately the same height for the two plots, while the spike at about (600,0) in Figure 1 is due to the fact that the butt does actually catch bullets. Here the template (in Figure 4) shows much less leakage, and our estimate for the probability of hitting the box at (3400,1400) is now 9.7×10^{-10}, which is 10^4 times smaller than for the range with a butt. This raises the question of whether stop-butts, which are supposed to make ranges safer, may in fact make them more dangerous. There are too many assumptions in our model which were not based on real evidence for this question to be answered here with any degree of confidence. We can say that this comparison does point up the sensitivity of the model to features of the terrain.

Risk in the stop butt shadow As well as the risk at the edge of the range, we also considered the risk near the middle line of the range, at a point behind the stop butt. It might be assumed that one is in a safe area in the shadow of the butt if one is close to the butt and the firing point is not visible. To test this we introduced a witness-box 10m square and 2m high, centred at (625,0). This witness-box is not visible from the point of fire. Our estimated hit probability per round per unit area is 3.2×10^{-6}. This may seem small, but if three people always stand in this box when firing occurs, then the expected number of hits on these people per million rounds is about 10. If a million seems a large number of rounds, the same figures imply that the probability of one hit in 10000 rounds is about 0.1. One conclusion that can be drawn from this is that it is not easy to guess from diagrams (or maps) how effective hills are going to be in protecting parts of the range from impacts, without doing a detailed calculation.

4 OTHER ASPECTS OF THE SOLUTION

4.1 Validation of solution

The complexity of our problem renders a direct analytic attack impossible. However, a simplified model which *is* analytically soluble and is not too simple for the results to be obvious, but which is not realistic in detail, has been developed to provide a check on the correctness of the computer codes. See Camfield *et al* (1995) and Pope (1995).

4.2 Error estimation

The validations referred to above provide evidence that the relevant computer codes compute correctly. The question still remains: can we determine how accurately the computation is done on any particular occasion? In fact we know how to do this. Because of the independence of the randomised systematic sampling, the total sample may be divided into independent *sections*, for each of which an estimate may be computed. The variability of the estimates between sections may be used to assess the variability of the simulation technique. A detailed account is in Lin and Pope (1995).

5 ACKNOWLEDGMENT

This work was supported under a research contract between the Australian Ordnance Council and the University of Newcastle.

6 CONCLUSION

Apart from the historical record, which is equivocal, we have at present inadequate information to form a reliable opinion about the standard of safety institutionalised in present procedures. The long-term objective of our modelling is to provide practical procedures which can be used after appropriate training by those responsible for the management of practices on weapons firing ranges. These procedures must be scientifically based and the level of risk that is involved in their use must be known to an appropriate degree of accuracy. For some weapons systems, this goal may be largely achievable soon, but much remains to do.

REFERENCES

B. Amies, M. W. Jarvis, A. L. Pope, R. Smith, and P. Young. Methodologies for the establishment of small arms safety criteria. Discussion paper, Central Studies Establishment, Department of Defence, Canberra, Australia, 1986. Prepared for the Australian Ordnance Council PPSC meeting of 6 March 1986.

ARDEC, Picatinny Arsenal. Physical properties of .50 caliber round. Private communication, 1990.

Australian Delegation. Briefing. Minutes of International Range Safety Conference, London, 16-18 May 1989.

S. Bentley. Aeroballistic data. Letter from Australian Army P&E Group, 26 November 1986a. Annex A to LOG(C)189-33-3.

S. Bentley. Various forms of the drag function and ballistic coefficient. Letter from Australian Army P&E Group, 13 January 1986b. Annex A to LOG(C)189-33-11.

John Camfield, Xun-Guo Lin, and Alun Pope. An exactly soluble range safety problem. *Journal of Ballistics*, 1995. To appear.

Xun-Guo Lin and Alun Pope. A new method of measuring the variabilty of estimated contours by simulation. These Proceedings, *MODSIM95*, 1995.

Alun Pope. A general method of calculating probability-based risk contours for firing-ranges. Working Paper RANSAF1, Central Studies Branch, Department of Defence, Canberra, Australia, 1988.

Alun Pope. Predicting the risk in the operation of weapons firing ranges. Newcastle Statistics Report 91-1, Department of Statistics, University of Newcastle, NSW Australia, 1991.

Alun Pope. Three-dimensional exactly soluble range safety problems. Newcastle Statistics Report 95-2, Department of Statistics, University of Newcastle, NSW Australia, 1995.

The Times, 22 April 1983.

The Times, 13 April 1988.

United Kingdom Ordnance Board. Weapon danger areas for small arms. Proceeding, 2 October 1984. OB 42321.

United Kingdom Ordnance Board. Danger areas for hill background ranges. Proceeding, 28 February 1989. OB 42577.

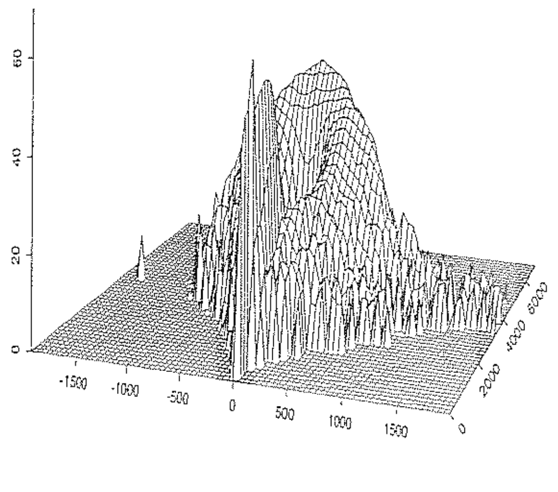

Figure 1: Perspective plot of logarithmic transformed probability density of final impact: range with butt.

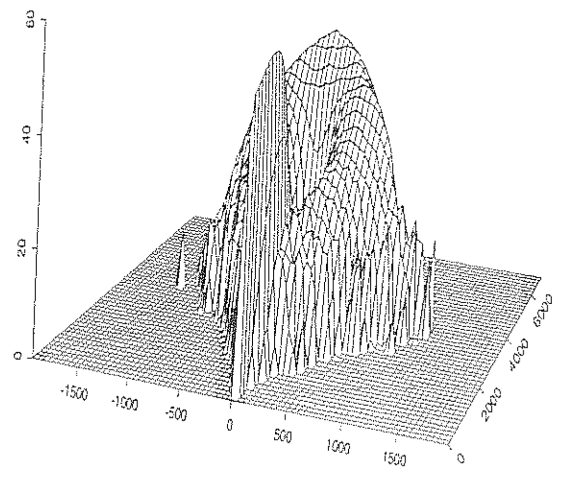

Figure 3: Perspective plot of logarithmic transformed probability density of final impact: flat range.

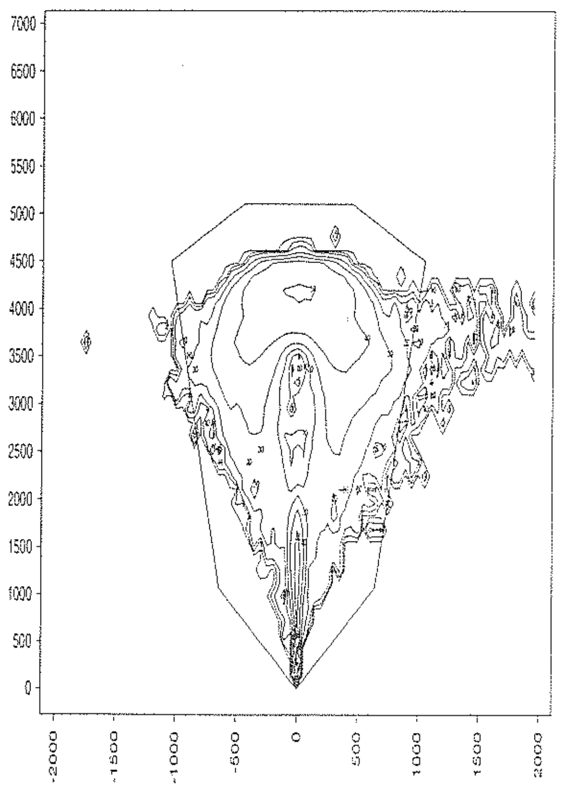

Figure 2: Contour plot of logarithmic transformed probability density of final impact: range with butt.

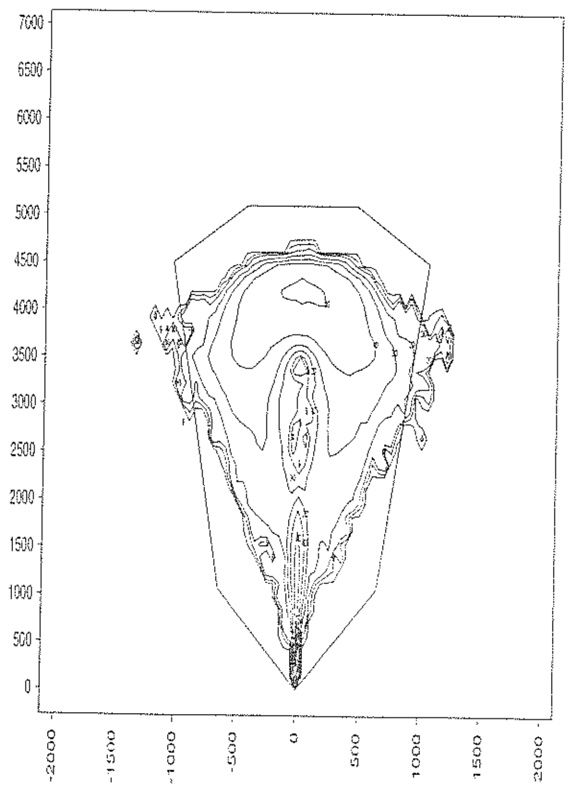

Figure 4: Contour plot of logarithmic transformed probability density of final impact: flat range.