

Development of a simplified parameterized model of firebrand landing from physics-based modelling

Matthew Zachary Kyng

Thesis submitted for the fulfilment of the requirement
for the degree of Doctor of Philosophy

Victoria University, Institute for Sustainable Industries & Liveable Cities (ISILC)

January 2025

Abstract

Bushfires pose a significant risk to lives and property, both within Australia and abroad. Firebrands (particles of hot or burning debris carried by air currents) have been identified as a particularly important aspect of how bushfires propagate and have been shown to cause a substantial proportion of structure fires at the wildland-urban interface. This research develops a parameterized model of firebrand landing distribution via analysis of simulated wildfires, using a physics-based model Fire Dynamics Simulator (FDS) which includes a low Mach number hydrodynamic model, a Large Eddy Simulation (LES) turbulence model, mixing controlled combustion model, heat transfer models, and a Lagrangian particle model.

FDS has been experimentally validated as a close approximation of real firebrand transport behaviour at laboratory scale. By simulating the distribution of firebrands at field-scale under different physical circumstances and flame dimensions, we aim to develop an appropriate statistical model of firebrand behaviour based on these parameters.

A diverse set of simulations under a variety of wind speeds, HRRs, flame dimensions (ratio of width to depth), and forest conditions have been completed and analysed. The results obtained from these simulations show that the shape of the distribution of firebrands throughout the domain varies significantly, with a distinct shift in the shape of the distribution occurring as the Width to Depth ratio (W:D ratio) of the fire increases.

Two distinct patterns of firebrand distribution have been observed from these simulations. In fires with a width to depth ratio of 2 or greater, a characteristic bimodal distribution is produced. A region of low firebrand landing density lies along the centre line of the domain, with a region of high landing density on either side. This two-peaked shape begins to form a very short distance away from the edge of the fire and persists for several hundred meters. The peaks gradually become less defined and sometimes can be seen to merge together at very long distances. In fires with a W:D ratio of 1, there is usually only one clearly defined peak.

A statistical model of firebrand landing has been developed that can capture this behaviour based on a set of probability density functions that can gradually shift between a single-peaked and a double-peaked shape. Trends in the parameters of the functions that describe the landing density in both the downwind and crosswind directions have been analysed. This model has been shown to produce a good quality of fit for the simulated data with a small number of input parameters. The physics of fluid motion and particle transport around the fire have also been investigated to provide an understanding of the model from a physics perspective.

Declaration

I, Matthew Kyng, declare that the PhD thesis entitled “Development of a simplified parameterized model of firebrand landing from physics-based modelling” is no more than 80,000 words in length including quotes and exclusive of tables, figures, appendices, bibliography, references and footnotes. This thesis contains no material that has been submitted previously, in whole or in part, for the award of any other academic degree or diploma. Except where otherwise indicated, this thesis is my own work.

I have conducted my research in alignment with the Australian Code for the Responsible Conduct of Research and Victoria University’s Higher Degree by Research Policy and Procedures.

Signature:



Date: 28/1/2025

Acknowledgement

I would like to thank my supervisors, Professor Khalid Moinuddin and Dr. James Hilton, for giving me the advice, feedback, and direction I needed throughout my research. Their feedback in proofreading my writing and working through the results of my simulations was invaluable for the progress of my research and the development of my skills in academic writing. Our frequent meetings kept me focused and on track with my research, where otherwise I may have become distracted. I would also like to give special thanks to Dr. Amila Wickramasinghe for helping me develop my skills with FDS, especially in the earliest days of my research at Victoria University. Without the examples given to me by Amila, I would not have become nearly as proficient with FDS as I needed to be in the early days of my research.

I would also like to thank my colleagues, Malavika Arun, Thattharani Kannangara, and Mohamed Sharaf for their good company while I performed my research at Victoria University. Our chats always brightened my day.

Finally, I would like to thank my friends and family, who supported me throughout my research. I am lucky to know so many people with a great depth of statistical, computational, and mathematical knowledge who have the patience to listen and understand while I describe obscure problems, even at odd hours of the day.

Summary of Chapters

This thesis consists of eight chapters. Each chapter deals with a separate topic or describes the results of a distinct set of simulation scenarios.

Chapter 1: Introduction

This chapter establishes the importance of this research in practical terms – such as the damage caused by firebrands in real fires, such as those at Duffy in Canberra, Australia, and introduces the reader to the general concept of the thesis.

Chapter 2: Literature Review

The Literature Review chapter examines the state of the field of firebrand generation, transport, and landing research. This serves to identify a gap in the research, which is then used to form the basis of a research question, as well as to provide context for the research undertaken in this thesis. This chapter goes into significant detail about specific relevant studies that form the basis for later choices, especially regarding the techniques of physics-based modelling and the properties of firebrands released from burning vegetation. Several studies relating to the transport and landing for firebrands are examined to identify existing trends and techniques that have been used to model firebrand distribution.

A review paper has been published with two other PhD students (attached as an appendix):

Wadhwani, R, K, Sullivan, C, Wickramasinghe, A, Kyng, M, Khan, N and Moinuddin, K. A review of firebrand studies on generation and transport. *Fire Safety Journal*, 2022, Vol 134, 103674.

Chapter 3: Model Description

This describes the mathematical basis of large eddy simulation as an approximation to Navier-Stokes equations, combustion and heat transfer equations, our choice of Haider & Levenspiel drag model, how FDS handles particle transport, and technical details of how the data is recorded and analysed.

Chapter 4: Scenario Development and Sensitivity Analysis

The Scenario Development and Sensitivity Analysis chapter describes the process of designing and testing the scenarios in which firebrand transport and landing are simulated. This chapter describes the reasoning behind choices relating to the physical and computational conditions of the scenarios – factors such as the dimensions of the domain and the boundary conditions

applied on each side of the domain, the intensity of the fire, atmospheric conditions, forest conditions, and firebrand characteristics. A sensitivity analysis is performed to demonstrate that a sufficiently fine LES grid size has been chosen for firebrand transport modelling.

Chapter 5: Simulation Results and Physics

This chapter describes the results of the simulations completed based on the previous scenario development chapter. The results of these simulations produce the dataset analysed in subsequent chapters. The physics of particle transport in the simulations is discussed, including the behaviour of the convective plume, vortices in the domain, and their impact on the trajectory of particles as well as the trajectory of particles as they move through the domain and how this results in the observed distribution of particles.

Chapter 6: Statistical Methodology and Model Development

The Statistical Methodology and Model Development chapter investigates the distribution of firebrands across the suite of scenarios that were successfully simulated and determines how each of the parameters impacts firebrand landing behaviour. This chapter begins with an exploratory data analysis that investigates basic trends in the data – e.g. how each parameter influences simple bulk properties of the distribution, such as average particle travel distance, coefficient of variation of particle distribution in the crosswind direction. Then, more refined quantitative methods are used to find statistical functions that can describe the distribution of particles. The distribution is investigated as a product of its marginal distributions and a coupling function, ultimately producing a function that takes in a set of parameters relating to the physical conditions of the scenario and producing a probability density function over a 2-D area.

Chapter 7: Final Model description and Performance Evaluation

This chapter presents a finalized version of the model developed in the previous chapter, describing the model succinctly and presenting it as a set of equations that can be used to produce an estimate of particle landing density at a given position downwind of a fire. The performance of the model is evaluated by comparing the model output against the simulated firebrand distributions obtained in Chapter 5: Simulation Results and Physics. The level of agreement or discrepancy between model results and simulated distribution is described statistically, and the performance of the model under different conditions is evaluated.

Chapter 8: Conclusion

Having developed a simplified parameterized model of firebrand landing distribution and evaluated its performance, this chapter discusses how the model addresses the research question and affirms a significant contribution to knowledge in the field of firebrand transport. Areas of further research are suggested to improve the model, as well as possible alternative statistical and computational techniques that may be used to develop future models.

Contents

Abstract.....	1
Declaration.....	2
Acknowledgement	3
Summary of Chapters	4
List of Figures	11
Chapter 1 Introduction	16
1.1 Research Context – Wildfires and Wildfire Modelling	17
1.2 Research Context – Firebrands	22
1.3 Research Aims and Firebrand Modelling	25
Chapter 2 Literature Review	26
2.1 Literature Review – Introduction.....	27
2.2 Firebrand Generation and Transport.....	27
2.2.1 Lab Scale Experiments	27
2.2.2 Field Scale Experiments	31
2.3 Models of Firebrand Transport and Distribution	35
2.3.1 Experimentally Validated Physics-based Models.....	35
2.3.2 Physics-based models	38
2.3.3 Empirical and Semi-empirical models.....	45
2.4 Research Gap	50
Chapter 3 Model Description.....	51
3.1 Introduction.....	52
3.2 Low Mach Number Approximation and Large Eddy Simulation in FDS	52
3.3 Governing Equations – Mass, Momentum, and Energy	53
3.3.1 Low Mach Number Approximation and Conservation of Energy.....	53
3.3.2 Mass Transport.....	54

3.3.3	Momentum Transport	55
3.4	Combustion Model.....	56
3.5	Lagrangian Particle Model.....	57
3.6	Solution Procedure – Time Step Selection	59
3.7	Solution Procedure – Predictor/Corrector.....	60
3.8	Summary	62
Chapter 4 Scenario Development and Sensitivity Analysis		64
4.1	Selection of scenarios	64
4.2	Thermo-physical properties of vegetation and firebrand characteristics.....	70
4.3	Domain size and boundary conditions	72
4.3.1	Domain Size.....	73
4.3.2	Inlet Conditions.....	77
4.3.3	Boundary Conditions	78
4.4	Grid Convergence Analysis	81
4.5	Discussion of Assumptions and Limitations.....	88
4.6	Summary	89
Chapter 5 Simulation Results and Physics.....		90
5.1	Introduction.....	91
5.2	Simulation Output.....	91
5.2.1	Data Output.....	91
5.3	Plume Characteristics.....	92
5.4	Vortex Behaviour.....	99
5.4.1	Wake Vortices.....	101
5.4.2	Counter-Rotating Vortex Pairs	106
5.5	Particle Trajectory.....	112
5.5.1	Stream Tracers	112
5.5.2	Downwind Dispersion and Initial Height	119

5.5.3	Crosswind Dispersion	126
5.6	Overall particle distribution and trajectory	132
Chapter 6 - Statistical Methodology and Model Development		140
6.1	Introduction.....	141
6.2	Data Visualization and Qualitative trends	142
6.2.1	Histograms	142
6.3	Kernel Density Estimation.....	145
6.4	Statistical Trends in Particle Distribution	146
6.4.1	Downwind Axis	146
6.4.2	Crosswind Axis.....	149
6.4.3	Trends across particle dimensions and geometry	151
6.5	Curve Fitting	155
6.5.1	Downwind Marginal Distribution.....	156
6.5.2	Crosswind Marginal Distribution	162
6.6	Describing the Joint Distribution	168
6.6.1	Copula Method.....	168
6.6.2	Conditional Probability Method	170
6.7	Model Fitting	177
Chapter 7 Model Description and Performance Evaluation		184
7.1	Introduction.....	185
7.2	Model Description	185
7.3	Downwind Marginal Distribution.....	185
7.4	Crosswind Conditional Probability Distribution	186
7.5	Model Performance.....	188
7.5.1	Performance Measurement Statistics	188
7.6	Overall Performance Statistics and Summary	194
7.7	Model Performance in Specific Scenarios	198

7.7.1	Open Field Cases	199
7.7.2	Forested Cases	205
7.8	Conclusion	210
Chapter 8 Summary and Conclusion		212
8.1	Introduction.....	212
8.2	Overview of Model Development	213
8.2.1	Physics-based Model Selection and Scenario Development	213
8.2.2	Simulation Results and Analysis of Physics.....	213
8.2.3	Statistical Methodology and Model Development	214
8.2.4	Model Performance.....	216
8.3	Recommendations for Future Research	217
References:.....		219
Chapter 9 - Appendix I		224
Chapter 10 - Appendix II.....		271
	Open Field Cases	271
	Forested Cases	283
Chapter 11 - Appendix III.....		296
	Open Field Cases	297
	Forested Cases	308

List of Figures

: TEMPERATURE ANOMALY MAP OF AUSTRALIA, JAN 1 – DEC 31 2019. [11]	18
: PLOT OF AUSTRALIAN MEAN TEMPERATURE ANOMALY. [12]	19
: PLOTS OF CALIFORNIAN HEATWAVE FREQUENCY, DURATION, AND INTENSITY, SHOWING A GRADUAL INCREASE OVER TIME SINCE 1950. [13]	19
: HISTOGRAM OF ERROR PERCENTAGE FOR A SELECTION OF WILDFIRE RATE-OF-SPREAD MODELS, FROM CRUZ AND ALEXANDER [17].	21
: PLOT OF FIREBRAND RELEASE RATE OVER TIME FOR SPECIES STUDIED IN MIGUEL ET. AL. [35]. EUCALYPTUS (EUC 01 AND EUC 02) IS SHOWN TO PRODUCE A RELATIVELY STEADY QUANTITY OF FIREBRANDS OVER TIME.....	30
: PLOT OF FIREBRAND LANDING DENSITY OBTAINED FROM PROJECT VESTA [37]. CONTOURS SHOW THE NUMBER OF FIREBRANDS LANDING PER SQUARE METER.....	33
: LOCATION OF FBP AND FCS SITES IN THE PINELANDS NATIONAL RESERVE, FROM [38].	34
: PLOT OF EXPERIMENTAL RESULTS OVERLAID WITH SIMULATED FIREBRAND DISTRIBUTION USING THE HAIDER & LEVENSPIEL DRAG MODEL.....	36
: DIAGRAM OF COMPUTATIONAL DOMAIN AND DISTRIBUTION OF FIREBRANDS (PARTICLE TYPE D1, 50KG/M ³ DENSITY).	39
: DIAGRAM OF THE INTEGRAL MODEL OF A CONVECTIVE PLUME WITH A GAUSSIAN PROFILE OF PLUME PROPERTIES.	42
: AFFINE RELATIONSHIP BETWEEN THE PARAMETERS OF THE LOG-NORMAL FUNCTION AND POWERS OF FLAME INTENSITY AND WIND SPEED. THESE PLOTS WERE GENERATED FROM SCENARIOS WHERE THE CHAR CONTENT OF PARTICLES WAS SET TO 39%.	44
: PLOT OF TERMINAL VELOCITY AGAINST PARAMETER Z.....	47
: DIAGRAM OF FORESTED AND NON-FORESTED DOMAINS, AND AN EXAMPLE OF FIREBRAND LANDING DENSITY IN EACH CASE (NOT TO SCALE).	65
: NEAR-SURFACE FUEL LOAD VS. AGE, FROM PROJECT VESTA [36].....	67
: FLAME TEMPERATURE VS. TIME OBSERVED IN PROJECT VESTA CONTROLLED BURNING EXPERIMENTS [36]	68
: IMAGE OF FIREBRANDS OBTAINED BY HOUSSAMI ET AL. [67]	71
: HISTOGRAM OF FIREBRAND PROJECTED AREAS FOR VARIOUS SPECIES – EUCALYPT RESULTS ARE EUC01 AND EUC02. FROM [35].	72
: MEASUREMENT INTERVALS BEFORE FOREST CANOPY	73
: CONVERGENCE OF AVERAGE WIND VELOCITY PROFILES IN RANGE 10-50M. CIRCLED POINT IS 10.3M/S AT 10M, WHERE U ₁₀ WAS SPECIFIED.....	74
: ILLUSTRATION OF DOMAIN WITH FOLIAGE AND DIMENSIONS.	75
: PLOT OF FIREBRAND LANDING DENSITY IN 54KM/HR WIND CONDITIONS, 4:1 ASPECT RATIO. EDGE OF FIRE LINE IS AT 400M.	76
: CHANGE IN WIND VELOCITY PROFILES AT 100M FROM THE INLET ACROSS 175M, 200M, AND 225M DOMAIN HEIGHTS.....	77
: COMPARISON OF FIREBRAND DISTRIBUTION HISTOGRAMS IN 46KM/HR WIND CONDITIONS. (A) DISTRIBUTION WITH OPEN SIDE BOUNDARIES AND (B) DISTRIBUTION WITH MIRROR BOUNDARIES.....	80
: OVERLAID VELOCITY PROFILES FOR 64KM/HR CASES, 150M FROM THE INLET.	84
: OVERLAID VELOCITY PROFILES FOR 64KM/HR CASES, AT THE CENTRE OF THE FOREST.	85
: DIAGRAM OF MESH LAYOUT THROUGHOUT THE DOMAIN, SHOWING COARSE AND FINE MESHES NEAR THE INLET.	87

: PLOT OF THE AVERAGE VERTICAL COMPONENT OF VELOCITY IN THE REGION IMMEDIATELY DOWNWIND OF THE FIRE. VELOCITY AVERAGED OVER 55 SECONDS DURING THE PERIOD OF PEAK HEAT RELEASE RATE. DASHED WHITE LINE REPRESENTS THE TOP OF THE FOREST IN THE FORESTED CASE.	94
: SIMULATED PLUMES WITH LITERATURE ZONATO MODEL OF PLUME RISE OVERLAID. DASHED WHITE LINE INDICATES FOREST HEIGHT IN THE FORESTED CASE.	97
: ZONATO MODEL OF PLUME RISE OVERLAID ON A SIMULATED PLUME FROM A SMALL, LOW-HRR FIRE CASE. THE PLUME RISES HIGHER THAN THE ZONATO MODEL WOULD PREDICT.	99
: DIAGRAM OF VORTICES FORMED BENEATH A CONVECTIVE PLUME IN A CROSSWIND. SOURCE: SHINOHARA AND MATSUSHIMA, [83]	101
: PAIRED WAKE VORTICES ON EITHER SIDE OF A FUEL BURNER HEAT SOURCE. THE VORTICES ARE ROTATING IN OPPOSITE DIRECTIONS – THE VORTEX ON THE LEFT SHOWING CYCLONIC ROTATION AND THE VORTEX ON THE RIGHT SHOWING ANTICYCLONIC ROTATION. SOURCE: CHURCH, SNOW AND DESSENS [82]	101
: EXAMPLE OF WAKE VORTICES MIGRATING DOWNWIND IN THE 54 KM/HR, 2:1WD, HIGH HRR, OPEN FIELD CASE. VORTEX DIRECTION IS INDICATED BY ARROWS. (A): VISUALIZATION AT 92.6S, (B): VISUALIZATION AT 102.6S.	103
: PLANES OF V-VELOCITY FOR OPEN FIELD AND FORESTED CASES (Z=30M). (A) DATA FROM THE 54 KM/HR, 2:1WD, HIGH HRR, OPEN FIELD CASE. (B) 54 KM/HR, 2:1WD, HIGH HRR, FORESTED CASE. (C) 54 KM/HR, 2:1WD, LOW HRR, OPEN FIELD CASE. (D) 54 KM/HR, 2:1WD, LOW HRR, FORESTED CASE. (E) 54 KM/HR, 4:1WD, HIGH HRR, OPEN FIELD CASE (F) 54 KM/HR, 4:1WD, HIGH HRR, FORESTED CASE	104
: COMPARISON OF PARTICLE DISTRIBUTION BETWEEN COMPARABLE FORESTED AND OPEN FIELD SCENARIOS USING KDE. THE FORESTED CASE SHOWS A MUCH MORE UNIFORM DISTRIBUTION OF PARTICLES IN THE CROSSWIND MARGINAL DISTRIBUTION.	105
: VISUALIZATION OF AVERAGE V-VELOCITY AND W-VELOCITY THROUGH THE SLICE X=500M FOR THE 46KM/HR, 2:1 WD, HIGH HRR CASES.	107
: PLOT OF FLAME EDGE Y-POSITION VS. VORTEX Y-POSITION	109
: PLOT OF Y_v/Y_f ACROSS Y_f AT X=500M. A TRENDLINE SHOWING A POWER LAW RELATIONSHIP IS SHOWN.	110
: KERNEL DENSITY ESTIMATE PLOTS OF PARTICLE LANDING DENSITY WITH VORTEX POSITIONS ALONG THE X = 500M LINE CIRCLED. .	111
: STREAM TRACERS FOR THE 46KM/HR, 2:1 W:D, HIGH HRR, FORESTED CASE. TOP: VIEW IN THE X-Z PLANE. BOTTOM: VIEW IN THE X-Y PLANE.	113
: STREAM TRACERS FOR THE 46KM/HR, 2:1 W:D, HIGH HRR, OPEN FIELD CASE. TOP: VIEW IN THE X-Z PLANE. BOTTOM: VIEW IN THE X-Y PLANE.	114
: STREAM TRACERS FOR THE 64KM/HR, 2:1 W:D, HIGH HRR, FORESTED CASE. TOP: VIEW IN THE X-Z PLANE. BOTTOM: VIEW IN THE X-Y PLANE.	115
: STREAM TRACERS FOR THE 64KM/HR, 2:1 W:D, HIGH HRR, OPEN FIELD CASE. TOP: VIEW IN THE X-Z PLANE. BOTTOM: VIEW IN THE X-Y PLANE.	116
: STREAM TRACERS FOR THE 54KM/HR, 4:1 W:D, LOW HRR, FORESTED CASE. TOP: VIEW IN THE X-Z PLANE. BOTTOM: VIEW IN THE X-Y PLANE.	117
: STREAM TRACERS FOR THE 54KM/HR, 4:1 W:D, LOW HRR, OPEN FIELD CASE. TOP: VIEW IN THE X-Z PLANE. BOTTOM: VIEW IN THE X-Y PLANE.	118
: PLOTS OF THE PROPORTION OF PARTICLES THAT LEAVE THE BURNING AREA FOR EACH CASE.	120

: GRAPH OF AVERAGE MAXIMUM PARTICLE ALTITUDE FOR PARTICLES THAT LEFT THE BURNING AREA FOR FORESTED AND OPEN FIELD CASES. BLUE BAR INDICATES AVERAGE MAXIMUM PARTICLE ALTITUDE, ERROR BARS SHOW 1 STANDARD DEVIATION AROUND THE AVERAGE.	122
: HISTOGRAMS OF PEAK PARTICLE ALTITUDE, DESCRIBING WHAT FRACTION OF PARTICLES ACHIEVE A GIVEN MAXIMUM ALTITUDE UNDER DIFFERENT CONDITIONS. DASHED LINES INDICATE TREETOP ALTITUDE.....	124
: PLOTS OF INITIAL PARTICLE ALTITUDE AGAINST FINAL PARTICLE LANDING POSITION IN THE X-AXIS FOR 37 KM/HR CASES. SAMPLE OF 100 PARTICLES FOR 42 PARTICLE TYPES, TOTAL OF 4200 PARTICLES IN EACH PLOT.	125
: PLOTS OF INITIAL PARTICLE ALTITUDE AGAINST FINAL PARTICLE LANDING POSITION IN THE X-AXIS FOR A SELECTION OF 46KM/HR AND 64KM/HR CASES. SAMPLE OF 100 PARTICLES FOR 42 PARTICLE TYPES, TOTAL OF 4200 PARTICLES IN EACH PLOT.	125
: PLOTS OF FIREBRAND INITIAL AND FINAL POSITIONS FOR A RANGE OF 4:1 WD, 54KM/HR CASES.	126
: PLOTS OF FIREBRAND INITIAL POSITION FOR 37KM/HR, HIGH HRR, OPEN FIELD CASES OF INCREASING W:D RATIO.	128
: PLOTS OF CHANGE IN PARTICLE Y-POSITION FOR DIFFERENT INITIAL POSITIONS FOR A RANGE OF 4:1 WD, 54KM/HR CASES.....	129
: KERNEL DENSITY ESTIMATE PLOT FOR THE 54KM/HR, 4:1 W:D, HIGH HRR, OPEN FIELD CASE.	130
: CHANGE IN PARTICLE Y-POSITION FOR A SELECTION OF NARROWER W:D FIRE CASES AT 54KM/HR AND 64KM/HR	131
: KERNEL DENSITY ESTIMATE PLOT OF PARTICLE DISTRIBUTION FOR THE 54KM/HR, 1:1 WD, HIGH HRR, FORESTED CASE. THIS CORRESPONDS TO THE CURVE SHOWN IN FIGURE 5-28B.	132
: PLOT OF PARTICLE TRAJECTORIES ON THE X-Y PLANE FOR A SELECTION OF TWO-PEAKED DISTRIBUTIONS. INITIAL PARTICLE POSITIONS ARE SHOWN IN RED, FINAL LANDING POSITIONS ARE SHOWN IN BLUE. GREY LINES SHOW THE PATH TAKEN BY EACH PARTICLE.	134
: COMPARISON OF PARTICLE TRAJECTORIES BETWEEN SIMILAR FORESTED AND OPEN FIELD SCENARIOS. THE FORESTED CASE PRODUCES A WIDER AND MORE EVEN DISTRIBUTION OF PARTICLES. INITIAL PARTICLE POSITIONS ARE SHOWN IN RED, FINAL LANDING POSITIONS ARE SHOWN IN BLUE. GREY LINES SHOW THE PATH TAKEN BY EACH PARTICLE.	136
: PLOT OF PARTICLE TRAJECTORIES ON THE X-Y PLANE FOR A SELECTION OF FORESTED CASES, SHOWING A WIDER AND MORE UNIFORM DISTRIBUTION OF PARTICLES. INITIAL PARTICLE POSITIONS ARE SHOWN IN RED, FINAL LANDING POSITIONS ARE SHOWN IN BLUE. GREY LINES SHOW THE PATH TAKEN BY EACH PARTICLE.....	137
: PLOT OF PARTICLE TRAJECTORIES ON THE X-Y PLANE FOR A SELECTION OF SINGLE-PEAKED OPEN FIELD CASES, SHOWING A LARGE DEGREE OF OVERLAP IN FINAL LANDING POSITION FOR PARTICLES THAT ORIGINATE ON OPPOSITE SIDES OF THE CENTER LINE. INITIAL PARTICLE POSITIONS ARE SHOWN IN RED, FINAL LANDING POSITIONS ARE SHOWN IN BLUE. GREY LINES SHOW THE PATH TAKEN BY EACH PARTICLE	139
: HISTOGRAM OF FIREBRAND LANDING FOR 46KM/HR, HIGH HRR, 1:1 WIDTH TO DEPTH, OPEN FIELD SCENARIO. THE COLOUR GRADIENT SHOWS THE LOG OF DENSITY.....	142
: HISTOGRAMS OF FIREBRAND LANDING FOR 46KM/HR, HIGH HRR, OPEN FIELD SCENARIOS. A) 1:1 W:D RATIO, B) 2:1 W:D RATIO, C) 4:1 W:D RATIO	143
: HISTOGRAMS OF FIREBRAND LANDING FOR 46KM/HR, 2:1 W:D RATIO, OPEN FIELD SCENARIOS. A) LOW HRR, B) HIGH HRR	144
: AN EXAMPLE OF A KERNEL DENSITY ESTIMATE PLOT WITH MARGINAL DISTRIBUTIONS. DATA IS FROM THE OPEN FIELD 46KM/HR, HIGH HRR, 4:1 WD CASE	145
: TRENDS IN AVERAGE DOWNWIND DISTANCE TRAVELLED.....	147
: TRENDS IN THE 10 TH DECILE OF DOWNWIND DISTANCE TRAVELLED.	148

: PLOT OF THE COEFFICIENT OF VARIATION IN THE X-AXIS FOR DIFFERENT SCENARIOS.	149
: VARIANCE OF PARTICLE DISTRIBUTION IN THE Y-AXIS, SHOWING LINEAR TRENDS WITH FLAME WIDTH	150
: PLOT OF THE COEFFICIENT OF VARIATION IN THE Y-AXIS.....	151
: <i>BOX PLOT OF PARTICLE DISTRIBUTION IN THE DOWNWIND DIRECTION, SEPARATED BY PARTICLE SPECIES. THIS PLOT SHOWS THE DISTRIBUTION OF EACH PARTICLE TYPE IN THE DOWNWIND DIRECTION — SHOWING THAT CERTAIN SPECIES AND GEOMETRIES ARE MORE PRONE TO PRODUCING LONG-TAILED DISTRIBUTIONS THAN OTHERS. IN PARTICULAR, LIGHTWEIGHT SPHERICAL AND CYLINDRICAL SPECIES PRODUCE MUCH LONGER TAILS THAN OTHER SPECIES.</i>	153
: BOX PLOT OF PARTICLE DISTRIBUTION IN THE CROSSWIND DIRECTION, SEPARATED BY PARTICLE SPECIES. THE LONG-TAILED DISTRIBUTION OF CYLINDRICAL PARTICLES CAN BE SEEN IN THE PLOTS OF PARTICLE IDS 24 THROUGH 30.	155
: KDE PLOT OF PARTICLE LANDING DISTANCE SEPARATED BY SPECIES. LIGHTER COLOUR INDICATES HIGHER PARTICLE ID NUMBER... 156	
: EXAMPLE HISTOGRAM AND KDE PLOT OF THE NATURAL LOG OF PARTICLE TRAVEL DISTANCE FOR AN EXAMPLE PARTICLE SPECIES.. 157	
: EXAMPLE PLOTS OF THE INVERSE GAUSSIAN CDF AND ECDF OF PARTICLE LANDING POSITION FOR TWO PARTICLE TYPES IN THE 64 KM/HR, 1:1 W:D, HIGH HRR, OPEN FIELD CASE. PARTICLE 41 HAS A CRAMÉR–VON MISES CRITERION OF 0.126, CONTRASTING WITH PARTICLE 24 WITH A CRITERION OF 3.31.....	159
: QUANTILE-QUANTILE PLOTS FOR THE INVERSE GAUSSIAN AND LOG T DISTRIBUTIONS. 64KM/HR 4:1 W:D HIGH HRR FORESTED CASE.	161
: KDE PLOTS ILLUSTRATING THE CHANGE IN DISTRIBUTION SHAPE ACROSS DIFFERENT W:D RATIOS FOR A SELECTION OF OPEN FIELD CASES.	163
: KDE PLOTS ILLUSTRATING THE CHANGE IN DISTRIBUTION SHAPE ACROSS DIFFERENT W:D RATIOS FOR A SELECTION OF FORESTED CASES.	163
: PLOTS OF A GAUSSIAN MIXTURE FIT TO THE CROSSWIND MARGINAL DISTRIBUTION SHOWING VARYING QUALITY OF FIT.TOP PLOT SHOWS A GOOD QUALITY OF FIT, THE BOTTOM PLOT SHOWS POOR QUALITY OF FIT.	165
: COMPARISON OF QUALITY OF FIT BETWEEN GAUSSIAN MIXTURE AND SKEWNORMAL MIXTURE FUNCTIONS.....	166
: PLOT OF A MODIFIED SKEWNORMAL MIXTURE FUNCTION FIT TO A SYMMETRICAL DATA SET.	167
: COMPARISON OF A KERNEL DENSITY ESTIMATE PLOT AND THE UNDERLYING COPULA. TOP: 64 KM/HR, W:D = 1, HIGH HRR, OPEN FIELD. MIDDLE: 46KM/HR, W:D = 2, HIGH HRR, OPEN FIELD. BOTTOM: 46KM/HR, W:D = 4, HIGH HRR, OPEN FIELD.	169
: DIVISION OF THE DOMAIN INTO SEGMENTS.	172
: PLOTS OF CONDITIONAL DISTRIBUTIONS FOR EACH SEGMENT. TITLE INDICATES THE DISTANCE AT WHICH THE SEGMENT BEGINS....	172
: PLOT OF THE SKEWNESS TERM OVER DISTANCE FOR FORESTED AND NON-FORESTED SCENARIOS. HYPERBOLIC CURVES FIT TO THE DATA SHOW SOME IMPACT OF FOREST ON SKEWNESS.	174
: HYPERBOLIC CURVES FIT TO SKEWNESS FOR CASES SEPARATED BY W:D RATIO. HIGHER W:D CASES HAVE CONSIDERABLY HIGHER SKEWNESS IN THE REGIONS NEAR THE FIRE.....	175
: DEPENDENCE OF THE LOCATION TERM AT X=400 ON REAL FLAME WIDTH.....	176
: CHANGE IN THE SHAPE OF SCALE OVER DISTANCE CURVES FOR DIFFERENT W:D RATIOS.	177
: EFFECT OF HRR ON μ PARAMETER OF THE INVERSE GAUSSIAN FUNCTION FOR DIFFERENT PARTICLE MASSES AND GEOMETRIES, SEPARATED BY FOREST CONDITION.	179
: EFFECT OF W:D RATIO ON μ PARAMETER OF THE INVERSE GAUSSIAN FUNCTION FOR DIFFERENT PARTICLE TYPES AND FOREST CONDITIONS.....	180

: LINEAR RELATIONSHIP BETWEEN μ AND $\ln(\text{MASS})$ FOR CYLINDRICAL PARTICLES.	181
: COMPARISON OF THE SIMULATED DISTRIBUTION OF PARTICLES (LEFT) AND THE STATISTICAL MODEL OF PARTICLE LANDING (RIGHT) FOR THE 46KM/HR, 4:1 WD, HIGH HRR OPEN CASE.....	182
: OVERLAID KERNEL DENSITY ESTIMATE PLOTS FOR THE 46KM/HR, 2:1 W:D, HIGH HRR, OPEN FIELD CASE.....	189
: OVERLAID KERNEL DENSITY ESTIMATE PLOTS FOR THE 37KM/HR, 1:1 W:D, LOW HRR, OPEN FIELD CASE. THIS CASE SHOWS AN EXAMPLE OF THE MODEL OVER-PREDICTING FIREBRAND DISTRIBUTION, ESPECIALLY IN THE FAR DOWNWIND AREAS.	190
: KERNEL DENSITY ESTIMATE PLOTS OF THE MARGINAL DISTRIBUTION IN THE X-AXIS FOR TWO CASES: (A) THE 46KM/HR, 4:1 W:D, HIGH HRR, OPEN FIELD CASE. (B) THE 37KM/HR, 1:1 W:D, LOW HRR, OPEN FIELD CASE.	191
: PLOTS OF MODEL PREDICTIONS FOR THE 46KM/HR, 2:1 W:D, HIGH HRR, OPEN FIELD CASE. (A) OVERLAID KDE FOR THE DOWNWIND REGION OF THE DOMAIN. (B) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE X-AXIS. (C) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE Y-AXIS. (D) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE X-AXIS (E) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE Y-AXIS.....	199
: PLOTS OF MODEL PREDICTIONS FOR THE 54KM/HR, 2:1 W:D, LOW HRR, OPEN FIELD CASE. (A) OVERLAID KDE FOR THE DOWNWIND REGION OF THE DOMAIN. (B) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE X-AXIS. (C) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE Y-AXIS. (D) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE X-AXIS (E) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE Y-AXIS.....	201
: PLOTS OF MODEL PREDICTIONS FOR THE 37KM/HR, 1:1 W:D, LOW HRR, OPEN FIELD CASE. (A) OVERLAID KDE FOR THE DOWNWIND REGION OF THE DOMAIN. (B) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE X-AXIS. (C) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE Y-AXIS. (D) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE X-AXIS (E) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE Y-AXIS.....	203
: PLOTS OF MODEL PREDICTIONS FOR THE 46KM/HR, 1:1 W:D, HIGH HRR, FORESTED CASE. (A) OVERLAID KDE FOR THE DOWNWIND REGION OF THE DOMAIN. (B) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE X-AXIS. (C) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE Y-AXIS. (D) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE X-AXIS (E) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE Y-AXIS.....	206
: PLOTS OF MODEL PREDICTIONS FOR THE 37KM/HR, 4:1 W:D, LOW HRR, FORESTED CASE. (A) OVERLAID KDE FOR THE DOWNWIND REGION OF THE DOMAIN. (B) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE X-AXIS. (C) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE Y-AXIS. (D) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE X-AXIS (E) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE Y-AXIS.	207
: PLOTS OF MODEL PREDICTIONS FOR THE 54KM/HR, 4:1 W:D, LOW HRR, FORESTED CASE. (A) OVERLAID KDE FOR THE DOWNWIND REGION OF THE DOMAIN. (B) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE X-AXIS. (C) OVERLAID KDE FOR THE MARGINAL DISTRIBUTION IN THE Y-AXIS. (D) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE X-AXIS (E) Q-Q PLOT COMPARING THE MODEL AND SIMULATED DISTRIBUTION IN THE Y-AXIS.	209

Chapter 1 Introduction

Contents

Research Context – Wildfires and Wildfire Modelling	17
Research Context – Firebrands	22
Research Aims and Firebrand Modelling	25

1.1 Research Context – Wildfires and Wildfire Modelling

Wildfires pose a significant risk to both lives and property, especially in communities that exist at the Wildland-Urban Interface (WUI)[1, 2]. When the conditions are right, any source of ignition could quickly lead to an out-of-control wildfire. High wind speeds, low rainfall, dry air, and a build-up of fuel over time can create the perfect conditions for a devastating wildfire. These conditions are likely to become more common in the future, as climate change increases global average temperature and alters rainfall patterns. Research by Jolly et. al. indicates that climate change has already led to a significant increase in the duration of fire seasons across the globe[3], and projections of future wildfire frequency by Huang et al [4] show a marked increase in the frequency of wildfires by 2050 due to climate change and altered land use patterns. Based on a range of Intergovernmental Panel on Climate Change (IPCC) emissions scenarios[5], these projections indicate an increase in wildfire frequency of 27% averaged across the globe, with certain areas (particularly Africa and Australia) potentially experiencing nearly 40% more wildfires in the summer months.

Climate change has already made high-risk wildfire conditions more common to a noticeable extent[6-8]. Recent wildfires in Australia, Europe, Africa, and the US have each destroyed thousands of homes, displaced tens of thousands of residents, and led to numerous deaths in communities at the wildland-urban interface. Of particular note among the recent wildfires was the 2019-2020 bushfire season in Australia and the 2020 California wildfires in the US. Both of these fire seasons were extremely destructive – the 2019-2020 Australian bushfires burned 24 million hectares, destroying 3000 homes and killing 33 people [9]. The 2020 California fires burned over a smaller area, but in a region with a much greater population density, resulting in comparable loss of life and destruction of property: 11,000 structures destroyed and 33 lives lost [10]. Detailed data on climate conditions in Australia and California during that period is readily available. The data for Australia in 2019 shows much of the continent experiencing a large amount of temperature anomaly – the difference between measured average temperatures in a given period and historical averages. This is shown graphically in Figure 1-1.

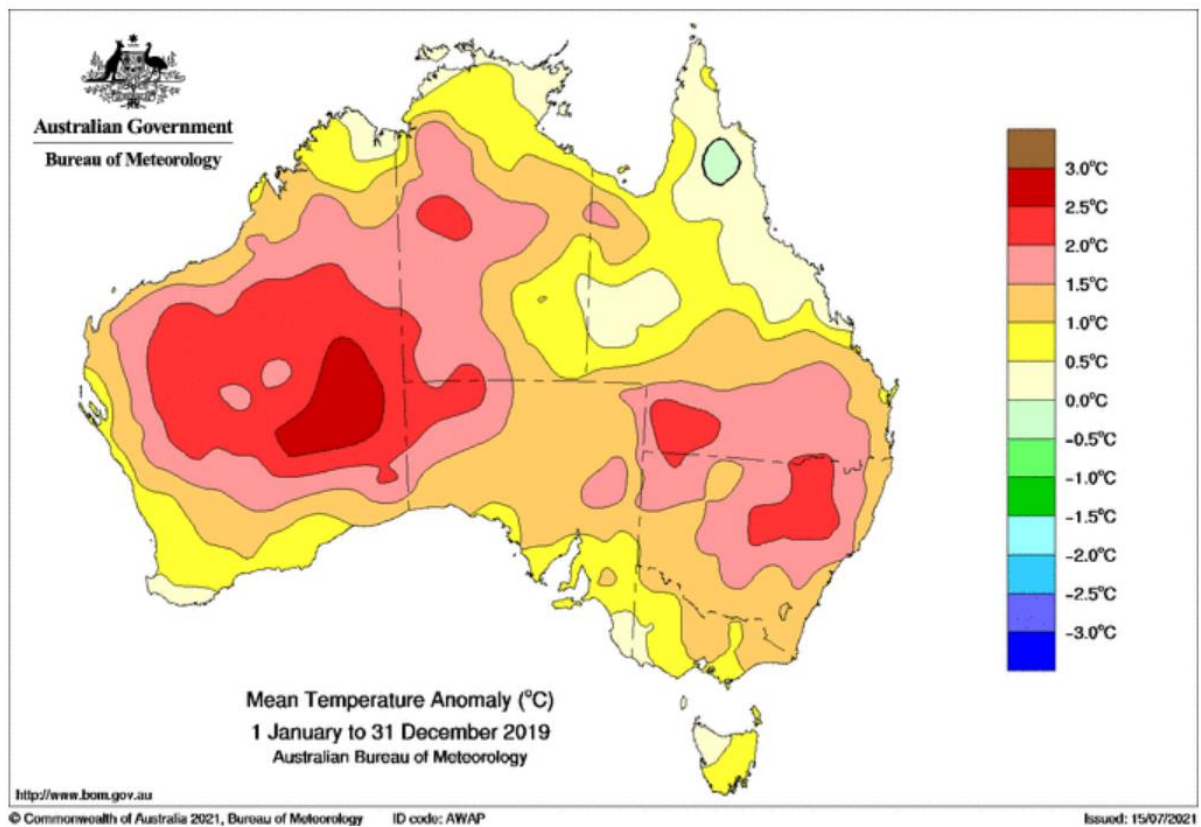


Figure 1-1: Temperature anomaly map of Australia, Jan 1 – Dec 31 2019. [11]

This is part of an ongoing trend in Australia – the average temperature anomaly each year is increasing, as shown in Figure 1-2. Combined with lower-than-average rainfall in many areas, the conditions in late 2019 led to the worst cumulative forest fire danger index (FFDI) values ever recorded. Similarly, data for California (Figure 1-3) shows an upward trend in the number of heatwaves, their duration, and their intensity over time.

Australian mean temperature anomaly




Figure 1-2: Plot of Australian mean temperature anomaly. [12]

More, Longer, and Hotter

Heatwaves affecting inland, urban California are on the rise




Figure 1-3: Plots of Californian heatwave frequency, duration, and intensity, showing a gradual increase over time since 1950. [13]

Given the significant and increasing danger wildfires pose to communities, there has been considerable effort for several decades to develop models of wildfire spread. These models typically aim to predict the rate of spread and intensity (or size) of a wildfire as it moves through

an area. These models can be divided into categories based on the approach taken to describe wildfire behaviour: empirical, semi-empirical, or physics-based.

Empirical models are developed by measuring the properties of real wildfires and controlled burns as functions of atmospheric, fuel, and terrain conditions. These are generally simple mathematical models with low computational requirements, designed to make predictions about fire behaviour at large scales (i.e. across tens to hundreds of kilometers). Often, these models make predictions in terms of the forward rate of spread (ROS) of a wildfire [14] based on easily measured parameters such as wind speed, ambient temperature, and rainfall conditions – such as the McArthur models of grassland and eucalypt forest fire spread[15].

Semi-empirical models combine theoretical principles with empirical data. Similar to empirical models, semi-empirical models are generally used to predict the behaviour of a fire at large scales. The Rothermel model [16] is an example of such a model, which is based on physical and theoretical principles that govern fire spread through vegetative fuels. It applies conservation of energy principles to calculate the rate of fire spread in various fuel models, taking into account the effects of fuel characteristics, slope, and wind speed. Despite its theoretical basis, the Rothermel model also relies heavily on empirical data to parameterize the equations used in the model. The model is used to predict how quickly a fire will spread in a specific type of fuel under given conditions. While these models can be useful in the role operational modelling of the fire propagation, modelling the landing pattern of firebrands released from these fires has been more difficult.

Physics-based models aim to describe the characteristics of a wildfire by approximating the physical processes that drive the wildfire – for example, simulating the fluid dynamics, heat transfer, and chemical processes that occur in real-world combustion. These models essentially simulate the interaction of fuel, fire, terrain, and atmosphere computationally, allowing for a highly customizable and accurate model of specific fire scenarios. As an example of the customizability of physics-based modelling, it is feasible to use a physics-based model to simulate the specific layout of a given room in a building for the purpose of modelling the effectiveness of a sprinkler system in fire suppression. However, these models quickly run into limitations in terms of computational resources – very fine spatial resolution is achieved at the expense of rapidly scaling requirements in terms of CPU operations and RAM requirements. As such, these models tend to be used for smaller-scale fire modelling applications.

Empirical and semi-empirical models are commonly used for large-scale fire behaviour modelling by emergency services. These models are well-suited to the requirements of emergency services operational planning, as they tend to rely on geographical information that can be gathered from surveys or remote sensing, such as terrain topology and vegetation characteristics. Empirical and semi-empirical wildfire models typically do not require large amounts of computation power to predict the path and speed of a fire. Developing such a model is, however, quite difficult – field-scale experiments are difficult and time-consuming to set up and require large areas of land to set aside for the experiment. These models vary in accuracy to a large degree. Cruz and Alexander [17] evaluated the performance of a selection of commonly used empirical wildfire rate-of-spread models against 1278 real wildfire rate-of-spread measurements under specified conditions. This evaluation revealed that most rate-of-spread models have large average errors when compared against real data – only a handful of models achieved an “exact” match (within 5%) of real observations at all, and the majority failed to produce an average result within 35% of the real observations. There was a trend across all conditions for models to under-predict the real rate of spread, as shown in Figure 1-4.




Figure 1-4: Histogram of error percentage for a selection of wildfire rate-of-spread models, from Cruz and Alexander [17].

Physics-based models, unlike empirical models, are generally computationally intensive and are very rarely used in an operational planning role. The high computational cost of running

detailed physics-based models of large areas means that these models are much slower than real-time. Simulating the combustion of a single tree in detail may take several hours or even days of computation on a current-generation CPU, and parallelization has diminishing returns past a certain point. Published literature on wildfire simulation illustrates this high computational cost – simulating a relatively small domain (50x30x25m) to model the effect of a grass fire on a model structure required between 16 and 71 hours of computation time on a 24-core node to complete just 20 seconds of simulation[18]. This makes the application of physics-based simulations of fire behaviour impractical for operational modelling. However, these physics-based models can be very accurate, and can simulate a wide range of scenarios. This makes these models ideal for simulating the behaviour of fires in complex situations well ahead of time. For example, Fire Dynamics Simulator (FDS) is a well-established tool for modelling fires, and has been used in the past to model building fires in detail[19]. By simulating combustion, pyrolysis, smoke generation, airflow, heat transfer, and particle motion, FDS can be applied to a variety of complex situations and determine how a fire may behave. A range of validation cases comparing FDS model outputs against experimental and benchmark model results are described in the FDS validation guide [20], including grassland fires and tree burning experiments. FDS has also been validated against laboratory-scale experiments for modelling firebrand particle transport by Wadhwani [21, 22].

1.2 Research Context – Firebrands

Firebrands are small particles of hot or burning debris released from burning vegetation or structures. These particles are carried away from the fire on air currents, lofted on convective plumes and carried downwind. These particles greatly complicate efforts to control wildfires in hazardous conditions – a large wildfire may release huge numbers of firebrand particles, and each one has the potential to ignite a spot fire some distance downwind of the fire front. When conditions are dry and large fuel loads are available, these spot fires may quickly grow to engulf a new area in flames or coalesce with the original fire front to effectively accelerate its spread. Spotting may also cause fires to spread beyond planned or naturally occurring firebreaks that could help contain a wildfire – igniting spot fires on the other sides of rivers or across clearings, for example.

In the 2003 fires at Duffy in the Australian Capital Territory, [23] firebrands were found to be the predominant cause of destroyed houses. Of the houses destroyed at Duffy, almost all of them were destroyed either by firebrands landing on the structure directly, or by a combination

of firebrands and radiant heat from nearby fires. Many of the destroyed houses were a considerable distance away from the fire front, where radiant heat from the original fire could not have been a major factor. This highlights the destructive potential of firebrand showers: even houses that were several streets away from the fire itself were at risk of destruction from firebrands transported on air currents.

Both destroyed and surviving houses with minor damage were studied to examine the mechanisms by which firebrands cause property damage. These structures reveal that certain architectural, landscaping, and design decisions make structures more prone to damage by firebrand attack. Due to the dry conditions around the ACT at the time of the fire, much of the vegetation in Duffy was dry, loose, and prone to ignition. Several destroyed houses were likely ignited by firebrands first starting spot fires that spread to nearby trees, which then produced more firebrands and exposed the structure to direct radiant heat.

Window, door, and fence designs that create re-entrant corners allow firebrands to accumulate in these areas, and these architectural elements are often made of combustible materials. Firebrands were also found to have entered houses through windows broken by exposure to radiant heat, or through external vents that were not adequately protected. Without swift intervention, spot fires ignited by firebrands landing in these locations can quickly destroy the whole structure.

Existing building standards in Australia relating to bushfire safety, such as AS3959 [24], do not currently have an adequate means to gauge the number of firebrands that are likely to land on a structure or to determine the likely characteristics of those firebrands. Understanding the characteristics, distribution, and accumulation patterns of firebrands on structures may allow for a significant improvement in building standards to resist firebrand attack.




Figure 1-5: Chart of houses damaged or destroyed at Duffy, categorized by the cause of damage.[23]

Firebrands can be categorized by the distance they travel from the fire front. In the CSIRO/AFAC Guide to Fire Spread Models by Cruz et al [25], three categories are described: short-range, medium-range, and long-range. Each of these categories is the result of a different kind of firebrand geometry and transport phenomena.

Short-range firebrands are those that land within 750m of the fire front. In general, these particles are not likely to have been lofted very high in the convective plume. Instead, they travel in relatively flat trajectories and are blown directly downwind of the fire. These particles are often observed in very large quantities in the areas immediately downwind of wildfires[26]. Because these firebrands travel only a short distance from origin to landing, short-ranged firebrands can have a wide range of sizes, shapes, and masses. Even particles that only burn or retain heat for a few tens of seconds can contribute to spotting over short ranges. Short-range firebrands have been identified as posing the greatest risk to structures at the wildland-urban interface, such as in the fires at Duffy.

Medium-range firebrands are those that land between 1000m and 5000m from the fire front. These firebrands are usually lofted in the convective plume for some distance before falling. Under certain conditions, large numbers of medium-range firebrands can land over an area in a short period of time, leading to a coalescence of spot fires a considerable distance from the original fire front. This can lead to complex patterns of airflow around the new fire front, causing particularly high intensity burning. [27]

Long range firebrands are those that travel over 5000m from the fire front. These firebrands are carried in the convective plume for a considerable amount of time. Relatively few firebrands manage to travel such long distances, and only specific types of particle (such as

long, slow-burning strips or tubes of bark) can remain hot enough to pose a significant chance of igniting a spot fire after being transported for so long. Because there are relatively few particles in this category, and they are distributed over a very long distance, spot fires ignited by these particles are unlikely to coalesce quickly with each other. Instead, they are likely to ignite small spot fires that grow over time to become a separate fire from the original fire front.

1.3 Research Aims and Firebrand Modelling

Given that efforts to control wildfires depend in part on the ability of operational planners to predict how a wildfire may spread, improving the speed and accuracy of operational models is desirable. The impact of short-range firebrands is of particular importance, given their potential to complicate firefighting efforts, ignite spot fires, and the large risk they pose to structures at the wildland-urban interface. Given the difficulty of capturing data on firebrand behavior in real wildfire conditions, numerical simulation may be a practical means to simulate the complex effects of the convective plume and vortices on the transport and distribution of firebrand particles for large-scale fires.

The primary aim of this research is to develop a simple, parameterized model of short-range firebrand distribution using physics-based modelling which can be used in operational modelling. This will be achieved by simulating firebrand transport and distribution in a range of scenarios with different atmospheric, fire, and forest conditions using an experimentally validated physics-based modelling approach in FDS. By simulating the distribution of firebrands under a range of different conditions, we may be able to analyze the results and develop a computationally inexpensive statistical model of firebrand distribution that could be integrated with existing operational models.

Chapter 2 Literature Review

Contents

1. Introduction.....	27
2. Firebrand Generation and Transport.....	27
2.1 Lab Scale Experiments.....	27
2.2 Field Scale Experiments.....	29
2.3 Models of Firebrand Transport and Distribution	35
2.3.1 Experimentally Validated Physics-based Models.....	35
2.3.2 Physics-based models	38
2.3.3 Empirical and Semi-empirical models.....	45
3. Research Gap	50

2.1 Literature Review – Introduction

Firebrand modelling is a complex field, with a significant body of existing research that must be considered before new developments can be made. There has been a considerable amount of research on particular aspects of firebrand behaviour – such as the generation of firebrands from vegetative sources, laboratory-scale studies of firebrand transport, wind-tunnel experiments investigating the forces acting on different firebrand geometries, and field-scale experiments investigating the distribution of firebrands from controlled burns. This body of existing research provides a basis from which a new parameterized physics-based model may be developed. Much of this literature review chapter is also presented in a published literature review paper [28] alongside other literature in related areas of firebrand and wildfire modelling.

2.2 Firebrand Generation and Transport

2.2.1 Lab Scale Experiments

The physical properties of firebrands depend primarily on the characteristics of the vegetation from which they originate and the conditions under which the firebrand was generated[29]. There has been considerable research in this field, trying to quantify the properties of firebrands generated by common plant species under a variety of physical conditions. It is necessary for any realistic physics-based model of firebrand distribution to have realistic particle geometries, masses, and physical properties, and to generate an appropriate quantity of particles for the conditions present. As such, this literature review will have a considerable focus on the particulars of firebrand generation from vegetative sources.

Manzello et. al. [30] describes a series of experiments that provide key insights into the rate of firebrand generation and characteristics of firebrands obtained from trees burning in the absence of wind. In this experiment, Douglas Fir trees of two different heights and diameters were burned in controlled conditions and the mass lost from each tree was recorded. A set of water-filled pans were placed around the tree in a fixed pattern, so that a sample of firebrands shed from the burning tree would be extinguished and captured. The physical characteristics of these firebrands were recorded, producing a detailed distribution of particle masses and geometries for each tree size.

From the results of this experiment, we can gather a number of useful parameters for our simulations – in particular, we obtain a mass and size distribution of firebrands from a common vegetative fuel source, as well as detailed information about the typical dimensions of

firebrands shed under no-wind conditions. It was found that almost all firebrands released from this vegetative source were cylindrical, typically long and narrow, with the majority of particles weighing less than 0.1g. The primary difference observed between the smaller and larger tree sizes was the relative proportion of particles with large masses, with very little difference in the characteristics of the firebrands otherwise.

Directly related to this experiment, Wickramasinghe et. al. [31] performed an inverse-analysis of the experiment by Manzello to determine the rate of firebrand generation from this vegetative source. This will be discussed at greater length later, in the section on physics-based modelling.

More complex experiments involving variable wind speed, different species of vegetation, and less consistent ignition sources (e.g. burning grass as an ignition source for a shrub) were conducted by Bahrani [32]. The focus of Bahrani's research is not only on firebrand generation and characterization of firebrands, but also on transport over short distances (14.3m) in laboratory conditions. This makes it a particularly valuable study – it is not only a source of information on the generation and physical characteristics of firebrands from a variety of species of vegetation, but it also includes a detailed description of the distribution of firebrands over a short distance.

The experiments described in Bahrani [32] investigate a wider variety of variables than the single-tree experiment described previously. In each experiment, a sample of vegetation from one of five species (Little Bluestem Grass, Chamise, Saw Palmetto, Loblolly Pine or Leyland Cypress) is placed in a test station, with an array of fans producing a specified wind speed (5.36 m/s, 11.17m/s, or 17.88 m/s) around the sample. For the higher wind speeds, an intermittent “gust” was introduced, increasing the wind speeds periodically. Downwind of the vegetation sample, an array of water-filled pans was placed at known distances to capture firebrands. Captured firebrands were weighed, photographed to obtain a projected area, and had their landing position recorded.

Of the several findings presented in this research, some of the most relevant are: the tendency of the experimentally obtained firebrands to follow a predictable relationship between mass and projected area for each species, the log-normal distribution of firebrand masses and areas, and the inconsistent relationship between wind speed and the quantity of firebrands generated. Notably, however, there is no clear choice of probability density function for firebrand travel distance that can be obtained from this experiment.

Further investigation of the effects of vegetation (species and fuel loading) was conducted by Hudson [33] at Oregon State University. Similar to the research by Bahrani, a series of burning experiments were conducted with a set of four common species of vegetation found in the US – Douglas Fir, Grand fir, Ponderosa Pine, and Western Juniper. Groups of one, three, or five trees were placed in a bed of straw, and a fan set up to produce a consistent downwind airflow across the bed of trees. The trees were then burned, and firebrands captured using a set of water-filled trays and fire-resistant fabric mats placed at fixed locations downwind of the trees. Char marks on the fabric mats were used to evaluate the amount of energy carried with each firebrand – hotter and longer-burning firebrands would tend to produce more intense charring. Digital image analysis software was used to automatically record the dimensions of both the collected firebrands and the char marks observed on the fabric sheets.

Analysis of the results obtained in these experiments reveals significant differences in firebrand geometry, generation rates, distribution, and tendency of firebrand to cause char marks between species. For example, on a mass-loss basis Douglas Fir was found to produce significantly more firebrands per kilogram than any other species, and Western Juniper produced the fewest – Douglas Fir producing firebrands at 1.7 times the rate of Western Juniper. However, the trend is reversed when it comes to the generation of char marks – Western Juniper produced the most char marks per kilogram of mass lost, nearly double that of Douglas Fir.

From these results, a series of regression models were produced, evaluating the importance of several factors in the generation of firebrands and char marks. In these models, the species of vegetation burned was found to be the single most significant factor for both ember generation and char marks. Based on these results, any attempt to model firebrand behaviour in a realistic forest fire must account for the particular characteristics of the species of vegetation that grow in the area, as even similar species (such as Grand fir and Douglas fir) may have significantly different firebrand generation characteristics.

Noting that each vegetation species produces firebrands with significantly different characteristics, it is worthwhile to consider some studies that investigate the properties of firebrands from Australian vegetation. There are several studies in this area, some studying certain species that produce firebrands with unique properties (such as research by Ellis [34] on the characteristics of Stringybark Eucalyptus firebrands) and others with a broader scope, or focus on more common species.

One study of particular interest that describes the firebrand generation rate and characteristics of an Australian species is Miguel et. al.[35], describing the generation rate and properties of firebrands from Southern Blue Gum (Eucalyptus Globulus). The experiment described in this article is somewhat similar to Manzello [30], in that a single sample of vegetation was burned without a source of wind to affect the generation or transport of firebrands. However, rather than capture firebrands with water-filled pans, this study used high-speed photography to capture images of firebrands as they were lofted from the burning sample. This allowed them to capture images of glowing firebrands in the convection column and record the times at which firebrands were released. Based on these images, they produced a series of curves of firebrand release rate over time and characterized the distribution of firebrand cross-sectional area for different species. This research found that Eucalyptus samples tended to release firebrands at a relatively steady rate during the burning period, unlike other species that quickly peaked and decreased. This is illustrated in Figure 2-1.

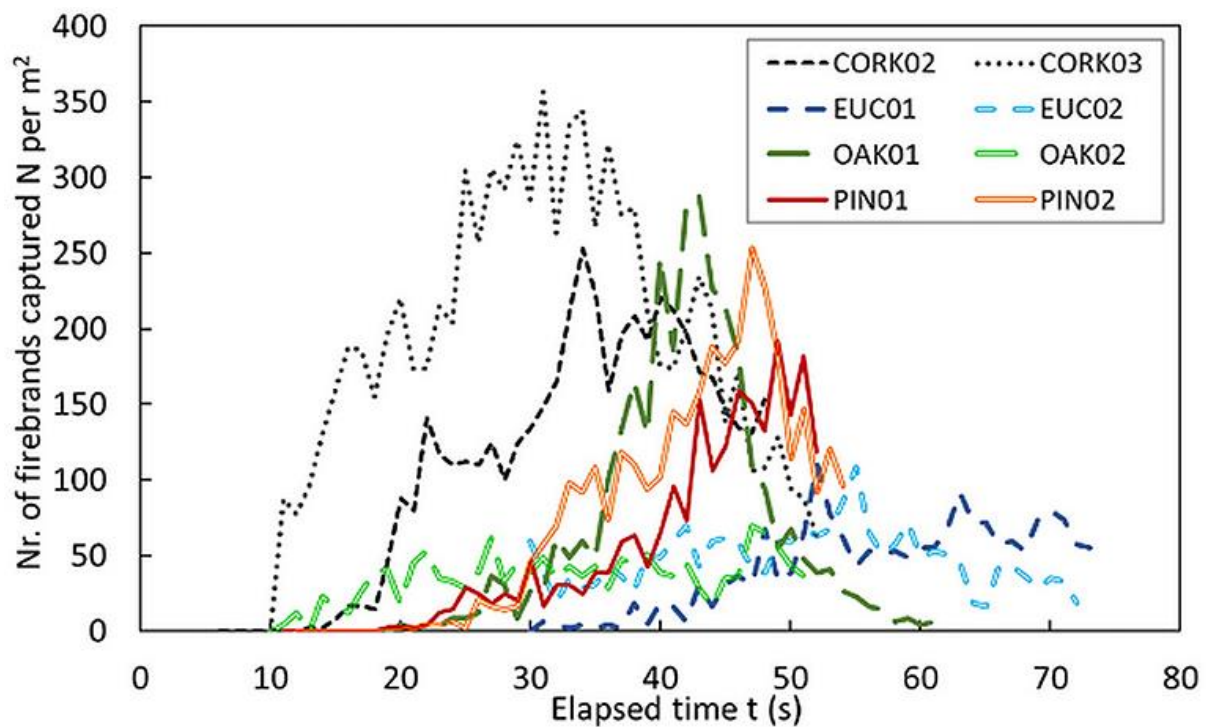


Figure 2-1: Plot of firebrand release rate over time for species studied in Miguel et. al. [35]. Eucalyptus (EUC 01 and EUC 02) is shown to produce a relatively steady quantity of firebrands over time.

This research may be highly relevant to the design of simulations of fires in eucalyptus forests – if the rate of firebrand generation does not change dramatically over time as fuel is consumed, it may be approximated as a constant release rate throughout the burning period.

2.2.2 Field Scale Experiments

While the physical characteristics of firebrands from specific species of vegetation can be researched at laboratory scale, it is difficult to determine the properties of a field-scale distribution of firebrands using laboratory-scale experiments. In addition to this, to properly simulate the generation and distribution of firebrands from a realistic wildfire, it is necessary to properly characterize the rate of spread and intensity of a realistic wildfire. As such, this section will focus on field-scale research on both firebrands and rate of spread models.

Project Vesta [36] was a CSIRO research project with the aim of developing an improved model of fire ROS in Eucalypt forests. Using an empirical approach, Project Vesta gathered data from several controlled burning experiments in Western Australia to characterize fire behaviour in relation to a number of key parameters. In addition to collecting data on fire ROS, Project Vesta also gathered a significant dataset concerning the distribution of firebrands downwind of the fire front.

The methodology of Project Vesta relied on thoroughly characterizing atmospheric and fuel conditions throughout each plot at the time of burning. For each of the 200x200m plots used in the experiment, a thorough measurement of fuel load was taken, measuring fuel load at five elevations: surface, near-surface, elevated, intermediate canopy, and overstorey canopy. Each of these layers was assigned two values – a “Percent Cover Score” (PCS) that represents fuel coverage as a proportion of area, and a “Fuel Hazard Score” (FHS) that represents characteristics of the fuel (e.g. proportion of dead plant matter, bark loosely or closely attached to trees). Each plot was then assigned an overall “Fuel Combustibility Score” produced by summing PCS*FHS across each layer. Fuel moisture content was measured in the surface and near-surface fuel layers, as well as in live samples taken from the “elevated” (shrub) layer.

Atmospheric conditions were recorded at multiple locations – both within the plot and at a weather station located at a fixed location clear of vegetation near the plot. Within the plot, a grid of anemometers at 5m elevation was used prior to burning to characterize the sub-canopy wind profile. Comparing data between points on this grid, it is possible to characterize the size, speed, and distance over which gusts of wind travel within the forest. Knowing the above-canopy wind speed, it is also possible to identify a correlation between above-canopy wind and the presence of gusts beneath the canopy. From this pre-burning atmospheric data, it was found that sub-canopy gusts do not travel a significant distance, losing speed quickly.

With these initial conditions thoroughly characterized, controlled burns were started. Flame ROS was measured using a series of nylon strings held under tension downwind of the fire – when the fire burned through the line, a data logger was triggered to record the time. Measurements of flame temperature were recorded at fixed locations using a series of thermocouples. In a selection of cases, firebrand distribution downwind of the 200x200m plot was measured using a series of plastic sheets laid out at fixed distances.

The results of this research are highly significant to the field of wildfire ROS and firebrand modelling. Using a series of regressions, an equation for predicting ROS in Eucalypt Forest was produced with an R^2 of 0.69:

$$R_{SS} = [R_t + b_0(U_{10} - U_t)^{b_1} e^{b_2 S_{fhs} + b_3 NS_{fhs} + b_4 NS_h}] \emptyset M_f \emptyset S_f \quad (2-1)$$

Where R_{SS} is the steady-state rate of spread, R_t is a known rate of spread at a threshold wind speed U_t , and U_{10} is wind speed at 10m in the open. S_{fhs} , NS_{fhs} , and NS_h are characteristics of the vegetation present in the area – the Surface fuel hazard score, Near-Surface fuel hazard score, and Near-Surface fuel height, respectively. The coefficients b_0 through b_4 are constants determined through a computational modelling process using S-PLUS. The terms $\emptyset M_f$ and $\emptyset S_f$ are correction functions that adjust rate of spread to account for fuel moisture content and terrain slope.

This equation was found to be in good agreement with historical data for fires of up to 2500m/hr spread, with the majority of historical fires in this category falling within 25% of predicted ROS.

In relation to firebrands, a trend in firebrand landing density was found, with two proposed equations that could describe landing density in the downwind and crosswind directions. The downwind landing density was approximated by:

$$D_{fb} = D_0 e^{-a_3 d} \quad (2-2)$$

Where D_{fb} is the number of firebrands landing per square meter at a distance d from the edge of the fire, D_0 is the number of firebrands landing at a point close to the fire, and a_3 is a constant.

Crosswind landing density was approximated by:

$$D_{fb} = D_{max} e^{-((d-d_{max})^2)/c_3} \quad (2-3)$$

As before, D_{fb} is the firebrand landing density at position d , this time measured across a transect perpendicular to the direction of the wind. D_{max} is the peak firebrand landing density measured across that line, which occurs at position d_{max} . The term c_3 is a constant.

The data collected also provides several useful insights about the shape of real-world firebrand landing distributions. The observed distribution produced curved landing density contours for the first 50m, followed by irregular distribution patterns further downwind. This is shown in Figure 2-2 .




Figure 2-2: Plot of firebrand landing density obtained from Project Vesta [37]. Contours show the number of firebrands landing per square meter.

Thomas et. al. [38] performed field-scale experimental research with a direct focus on firebrand generation and distribution, in contrast to Project Vesta, which was primarily focused on rate of spread modelling with firebrand research as a secondary concern. An experiment was conducted on a 28ha section of woodland in the Pinelands National Reserve in New Jersey, a site with a temperate climate and relatively flat terrain. The vegetation present in the canopy layer was predominantly pine trees, with a wider variety of species in the sub-canopy layer.

Detailed meteorological data and video footage was collected at three recording sites inside the plot to be burned, referred to as Fire Behaviour Packages (FBP) X, Y, and Z. Three firebrand Collection Sites (FCS) were set up with an array of water-filled aluminium cans just outside of the plot to be burned. Video cameras were set up to record the cans, so that the quantity of firebrands captured over time could be recorded and correlated with flame characteristics recorded at FBP sites. The FCS locations are labelled X, Y, and Z based on their proximity to the FBP sites, as shown in Figure 2-3.




Figure 2-3: Location of FBP and FCS sites in the Pinelands National Reserve, from [38].

The overall distribution of firebrand size and mass obtained in this experiment is similar to the distributions obtained in other experiments involving similar species – such as those by Manzello at both lab-scale [30] and field-scale. However, the distribution of particles between each FCS location was not uniform. As the fire propagated from FBP site Z to sites Y and X, the intensity of the fire reduced. The relative intensity of the flame at each FBP was found to dramatically alter the proportion of captured firebrands in each size category at each FCS.

Firebrands were divided by size into six categories, from less than $5 \times 10^{-5} \text{ m}^2$ to greater than $5 \times 10^{-4} \text{ m}^2$. FBP site Z experienced the highest flame intensity, estimated to be $12.59 \pm 5.87 \text{ MW/m}$. This was approximately 170% of the intensity at FBP site X. The firebrand flux for each particle size category at FCS Z was considerably higher than that at FCS X and Y, indicating a strong relationship between fire intensity and firebrand flux. The effect may be

more pronounced for larger categories of firebrands, but the small total number of firebrands captured in these large size categories made the results highly variable.

2.3 Models of Firebrand Transport and Distribution

2.3.1 Experimentally Validated Physics-based Models

Research into firebrand modelling using both experimental and physics-based approaches has produced several valuable results that this thesis will build upon. The validation of physics-based modelling of firebrand transport by experimental research is fundamental to the methodology of this project.

Research by Wadhwani et. al. [39] describes a process of experimental validation of a physics-based model of particle transport, as well as a thorough process of computational modelling of the thermal properties of firebrands and their distribution over an area. This study makes several key findings and is of significant importance to the methodology of this project.

Wadhwani conducted a series of lab-scale experiments with an artificial firebrand projector, capable of launching both burning and non-burning model firebrands over an area. The velocity profile of air in the firebrand generator was carefully controlled to ensure consistency between runs. Particle exit velocity from the generator was determined by analysis of video recordings using specialized software – measuring the distance each firebrand travelled between frames. The final landing distribution of the firebrands was also recorded, producing a series of contours describing the observed landing density of firebrands. Three different firebrand geometries were used in these experiments – cylindrical, cubiform, and square disc shapes.

Once the experimental data from the firebrand projector experiments was obtained, a series of computational simulations in FDS were set up with the same initial conditions. Each simulation implemented a different drag model for the Lagrangian particle model in FDS, testing a range of well-established models from literature. Of these models, the best performing model was the Haider & Levenspiel drag model [39]. This model consistently produced simulated results that closely overlapped experimentally obtained data at all wind speeds tested. Figure 2-4 shows the degree of overlap for the highest wind speed, with non-burning cubiform particles:




Figure 2-4: Plot of experimental results overlaid with simulated firebrand distribution using the Haider & Levenspiel drag model.

The Haider & Levenspiel drag model [40] calculates drag coefficient based on the sphericity of the particle, the Reynold's number, and a set of empirical correlations. The equations for this drag model are presented in Chapter 3: (3-20)through (3-24).

Following from this validation of the FDS Lagrangian particle model with Haider & Levenspiel drag, a simulation of firebrand transport at field-scale was conducted. Wadhwani [39] simulated a large section of forest in FDS, with a computational domain 1000x160x80m in size and a grid size of 2x2x1m. Firebrands were introduced along a plane within the forest while a heat source of a known intensity produced a convective plume. The landing positions of the firebrands were recorded, separated by particle geometry and size. The distribution of firebrands over the domain was observed to be qualitatively similar to distributions recorded in field-scale studies of firebrand behaviour, such as Project Vesta and Thomas et al. [38] Particle geometry was found to play an important role in the shape of the distribution – particles with lower sphericity generally dispersed over a wider area and had a greater average travel distance.

Tohidi [41] [42] describes a series of experiments and statistical analyses conducted with the aim of developing and validating a numerical model that can describe the motion of firebrands without certain simplifying assumptions used in previous models. By comparing and validating a variety of numerical methodologies used for particle transport, plume development, and turbulence modelling against experimental results, Tohidi produces a highly accurate numerical model of particle transport, but it is computationally very expensive.

The analysis of firebrand generation presented in Tohidi [42] is in part based on work by Manzello, including the single tree burning experiments described previously [30]. Data from Manzello's single tree burning experiments were combined with other results from literature to produce a model of the stresses acting on firebrands before they break off from a tree. Combined with physics-based models of tree geometry and the forces generated by convective airflow around a burning tree, Tohidi presents a model for determining the quantity and characteristics of firebrands that may break off from a tree. Combined with observed results from literature, a number of points can be taken from this analysis:

- Most firebrands released from a single tree burn have a large enough aspect ratio of length to width that they were likely broken off from the tree by the force of convective airflow around the tree.
- Firebrand surface area is proportional to mass to the $2/3$ power.
- Surface area and aspect ratio are not strongly correlated.

Following this analysis, a series of experiments were conducted to investigate the performance of a 6-degree-of-freedom numerical model in describing the distribution of firebrands landing after a period of freefall. Model firebrands of varying aspect ratio, size, and density were dropped from a known height in the absence of wind, and their landing distribution recorded via an image analysis system that captured the firebrand's position at the exact moment the particle landed. The experimentally obtained particle landing distribution was then compared with particle landing distributions predicted by the 6- degree of freedom (DOF) numerical model in a series of statistical tests. The numerical model was found to be highly accurate in all cases.

The behaviour of bent-over plumes in a cross flow (such as those produced by wildfires with ambient wind conditions) is the next area of investigation. This focuses on the evaluation of a number of existing models in literature – comparing the predictions of each with results from a number of plumes generated in previous experiments. This part finds that wildfire plumes are

heavily influenced by changes in wind speed with altitude, as well as defining key ratios of plume size and momentum for which different modelling approaches are appropriate.

Finally, a series of wind-tunnel experiments were conducted using a set of non-burning polyurethane model firebrands [41]. Model firebrands of three different aspect ratios (1, 4, and 6) were released from a fixed starting position in a specifically prepared wind tunnel. Conditions in the wind tunnel were controlled to produce three distinct centre-line wind velocities, as well as three distinct wind speeds at a fixed height of 0.04m above the floor of the wind tunnel (boundary layer velocity) – producing nine distinct configurations of wind speeds in the tunnel with different velocity profiles.

Upon releasing the model firebrands in the wind tunnel, a digital camera and specialized software recorded the exact trajectory of each model firebrand released. This produced a series of firebrand landing probability density functions (PDFs), which were then analysed. The clearest trend in all firebrand trajectories was that the final maximum downwind distance was strongly related to maximum lofting height – regardless of firebrand aspect ratio, the PDFs of the ratio of distance travelled to maximum lofting height had significant overlap.

2.3.2 Physics-based models

Wang [26] used a physics-based approach to derive a simple but useful mathematical description of the mechanism of transport for firebrands in a convective plume. By making certain simplifying assumptions about the geometry of firebrands, Wang derives an expression for the maximum radius of a spherical firebrand that can be lofted by a given plume velocity. This expression is then combined with results from literature to develop several further mathematical relationships that characterize firebrand behavior – such as the maximum lofting height, emission rate of firebrands as a function of firebrand radius, and the effect of reduced flame depth near the edge of a forest.

Using these relations and experimental results obtained from experiments performed by Hage [43] and Tarifa [44], a Rayleigh distribution is chosen to model the frequency of firebrands of a given size landing a distance x from the fire front. Combined with the function of firebrand emission rate as a function of radius developed previously, Wang presented a model of firebrand landing based primarily on simple theoretical principles.

To test this model, Wang compared the model's predictions to observed firebrand landing distributions from the 2003 fires at Duffy[45]. While the overall tendency of firebrand landing

density to decrease with distance agreed with observations of burned houses at Duffy, the model tended to significantly over-predict the chance of property fires in close proximity to the fire (less than 200m) and under-predict at longer ranges (over 300m).

Himoto and Tanaka [46] developed a simplified mathematical model of firebrand transport and landing for disk-shaped firebrands in a turbulent boundary layer based on computational simulation of fluid motion around a fire with LES, and simulation of particle transport with a physics-based approach. Their approach begins with a description of the forces and moments acting on a disk-shaped particle, describing the forces of drag, lift, and gravity on the particle, as well as a moment \mathbf{M} acting on the angular momentum of the particle. They approximate the drag coefficient, lift coefficient, and coefficient of the moment \mathbf{M} of the particle at different angles to the wind using relatively simple mathematical approximations of experimental results.

Using these approximations, they simulated the transport of disk-shaped particles with a range of different dimensions and densities in a clearly defined scenario. Wind was introduced into the domain at an inflow boundary with turbulence calculated in advance and a defined average velocity profile. The sides and top of the domain were set to a free-slip boundary condition. A square heat source of side length $D=0.08\text{m}$ and heat release rate (HRR) 4.0kW was simulated at ground level. Particles were inserted at height D above the downwind edge of the heat source. The domain of the simulation is illustrated in Figure 2-5:




Figure 2-5: Diagram of computational domain and distribution of firebrands (Particle type D1, 50kg/m^3 density).

To describe particle transport in this scenario, they first derive an expression for the maximum vertical height of a particle as it is lofted by the convective plume and carried forward by the wind. This expression is based purely on the balance of forces acting on the particle, rather than a simulation result – setting the force of buoyancy from the convective plume equal to the force of gravity. This expression is:

$$Z_p \cong \left(\frac{\dot{Q}}{\rho_\infty C_p T_\infty g^{\frac{1}{2}} d_p^{\frac{5}{2}}} \right) \left(\frac{\rho_\infty}{\rho_p} \right) \cdot d_p \quad (2-4)$$

Where ρ_∞ and T_∞ are the density and temperature of air under ambient conditions, \dot{Q} is the heat release rate, C_p is the heat capacity of air, g is gravitational acceleration, and ρ_p and d_p are the density and width of the particle. After reaching the maximum height Z_p , they assume that the dominant force acting on the particle in the vertical direction will be gravity. This means that the time taken for the particle to fall from height Z_p can be approximated as the time taken for a particle in freefall from that height to land on the ground. It is also assumed that after reaching height Z_p that the particle will be travelling horizontally at a speed equal to the speed of the wind. This allows for a convenient expression for x'_p - the horizontal distance travelled after reaching maximum height Z_p :

$$x'_p = \sqrt{2} \left\{ \frac{U_\infty}{(g d_p)^{\frac{1}{2}}} \left(\frac{\rho_\infty}{\rho_p} \right)^{-\frac{3}{4}} \left(\frac{\dot{Q}}{\rho_\infty C_p T_\infty g^{\frac{1}{2}} d_p^{\frac{5}{2}}} \right)^{\frac{1}{2}} \right\} \cdot d_p \quad (2-5)$$

Since the majority of the distance travelled by the particle takes place after reaching maximum vertical height, it is expected that the expression for total distance travelled x_p will be closely related to x'_p . They then define a dimensionless parameter B^* by normalizing the equation for x'_p by the characteristic length D , and describe overall travel distance x_p in terms of B^* and D :

$$B^* = \frac{U_\infty}{(g D)^{\frac{1}{2}}} \left(\frac{\rho_\infty}{\rho_p} \right)^{-\frac{3}{4}} \left(\frac{d_p}{D} \right)^{-\frac{3}{4}} \left(\frac{\dot{Q}}{\rho_\infty C_p T_\infty g^{\frac{1}{2}} D^{\frac{5}{2}}} \right)^{\frac{1}{2}} \quad (2-6)$$

They then seek to explain the shape of the distribution of particles obtained from the simulation in terms of this parameter B^* . A log-normal distribution was found to fit the distribution of

particles in the downwind direction quite well, with parameters μ_x and σ_x . It was observed that these parameters can closely approximated as functions of B^* :

$$\frac{\mu_x}{D} = 0.47B^{*2/3} \quad (2-7)$$

$$\frac{\sigma_x}{D} = 0.88B^{*1/3} \quad (2-8)$$

In the crosswind distribution, a normal distribution with parameters μ_y and σ_y was fitted to the data obtained from the simulation. In contrast to the downwind distribution, the parameters μ_y and σ_y were not found to depend on B^* , instead being well-approximated as a constant:

$$\frac{\mu_y}{D} = 0.00 \quad (2-9)$$

$$\frac{\sigma_y}{D} = 0.92 \quad (2-10)$$

These equations and distributions combine to produce a convenient and simple model of particle landing density for square disk-shaped firebrands with good performance in many cases. However, the model does seem to have some limitations – in particular, it seems that the relationships described by this study perform best for very low-density particles. As particle density increases, the log-normal function with parameters chosen by this model becomes a less accurate approximation of the simulated distribution. The discrepancy between the fitted distribution and the simulated data is noticeable for particles of 150kg/m^3 . This may indicate that the model needs some additional terms to account for particle density, and that certain assumptions about the acceleration of particles in this model may not hold for some real disk-shaped firebrands. This model also lacks a description of the joint distribution of the particles – characterizing the distribution of particles in each direction separately is not sufficient to reproduce the observed shape of the joint distribution.

Sardoy et. al. [47, 48] developed a model of firebrand distribution from a line fire using a physics-based approach. Rather than simulate a fire and convective plume with a computational fluids dynamics (CFD) approach, this model uses an integral approach to approximate plume behaviour with a Gaussian profile of plume characteristics. This produces a simple and

computationally inexpensive approximation of the velocity and temperature properties of the plume, as illustrated in Figure 2-6 .




Figure 2-6: Diagram of the integral model of a convective plume with a Gaussian profile of plume properties.

As disk-shaped firebrands have been found to travel the furthest and pose the highest risk of spotfire ignition, the particles investigated by this study are all of square disk geometry. Particle motion is calculated by solving equations governing the kinematic and angular momentum of the particle as it moves through the velocity field. In addition to this, the pyrolysis and combustion of particles during transport is modelled in detail, using a model described in another study by Sardoy [48]. This model assumes that particles retain the same shape (flat disk) but lose mass and reduce in size along one dimension (diameter) as combustion occurs.

Twenty scenarios with six different wind speeds and four flame intensities were computed. Flame intensities ranged from 10 to 40 MW m⁻¹, and wind speeds ranged from 6.70 to 17.88 m s⁻¹. For each of these simulations, vertical wind speed at the center of the plume at canopy height (10m above ground level) was taken to be a constant 12 m s⁻¹ based on previously obtained results from CFD simulation [48]. The diameter of the plume varied depending on flame intensity.

10,000 disk-shaped firebrands were released into this pre-computed plume and wind field at random throughout the diameter of the plume at the top of the canopy. The diameter, thickness, density, incident angle, and initial velocity of the particles were chosen at random within a

range of potential values. Densities ranged from 50 to 300 kg m⁻³, thickness from 1mm to 3mm, and diameter ranged from 4cm to 10cm.

While Sardoy obtained several useful results from this simulation relating to particle mass at time of landing, the primary result we are interested in is the shape of the distribution for short-range particles. As in the previously discussed study by Himoto and Tanaka [46], Sardoy found that the log-normal distribution had a good quality of fit for the simulated distribution of particles in the downwind direction. However, they found a different relationship between the conditions of the simulation and the values of μ_x and σ_x . Rather than using the dimensionless parameter B^* , Sardoy found that μ_x and σ_x follow an affine relationship with powers of flame intensity and wind speed. The parameters of this affine relationship and the powers of flame intensity and wind speed depend on the Froude number of the plume. The quality of fit for this relationship is very good, as shown in Figure 2-7 .

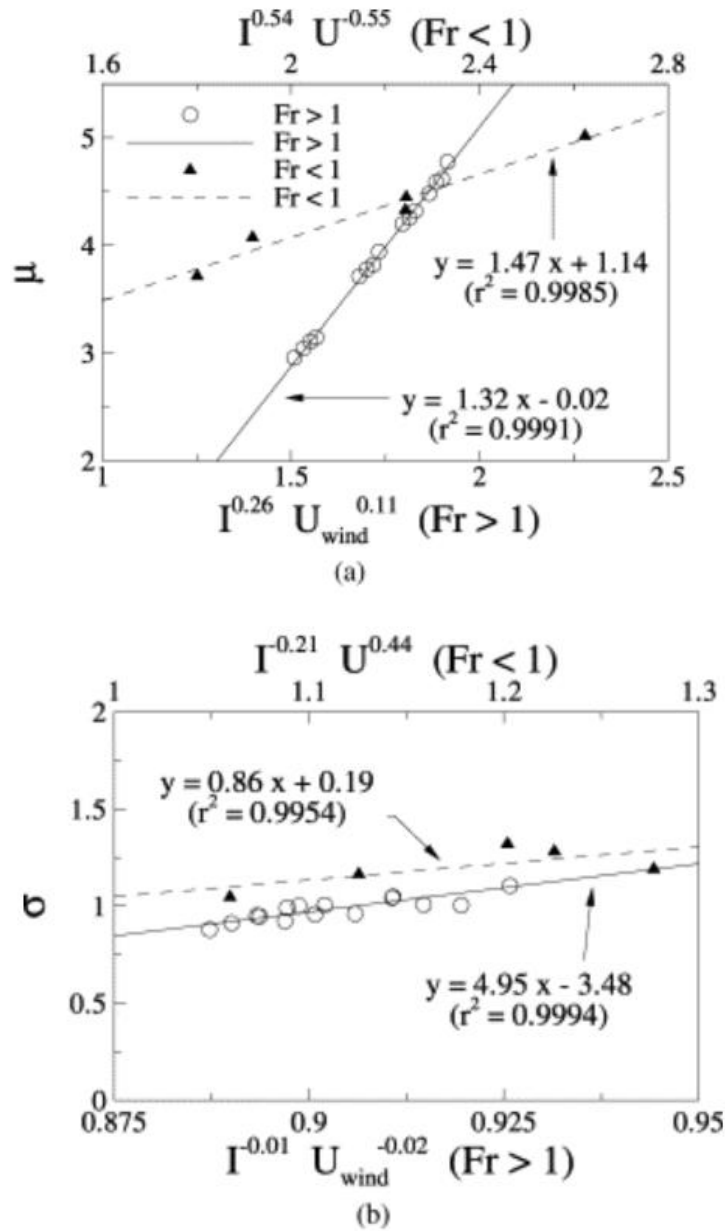


Figure 2-7: Affine relationship between the parameters of the log-normal function and powers of flame intensity and wind speed. These plots were generated from scenarios where the char content of particles was set to 39%.

This model of the distribution of short-range firebrands has some useful qualities that previously described models lack. This model describes the distribution of a population of firebrands with a range of randomized characteristics, rather than one particular type of firebrand with specific characteristics. The relationship used to predict μ_x and σ_x depend on some relatively easily observed characteristics of the fire and atmospheric conditions – the intensity of the fire and wind speed. These characteristics make this model more broadly applicable than some of the other models discussed in this chapter.

2.3.3 Empirical and Semi-empirical models

Many existing models of firebrand behaviour were developed using purely empirical or semi-empirical techniques. Some of these models do not aim to produce a function that predicts the landing density of firebrands at a given distance, but were developed to predict the maximum distance a spot fire could be ignited. Although these models were developed with a different goal in mind, the underlying experimental research and physics that were used to develop these models are relevant to the development of a practical simplified model of firebrand distribution. Many of these models were developed using simple approximations of the underlying physics of firebrand transport, with a limited range of inputs and no computationally expensive calculations required.

McArthur [15, 49] developed a model of maximum spotting distance for use alongside the Mk. 5 Forest Fire Danger Meter (also referred to as the Mk. 5 Forest Fire Danger Index, or FFDI) for use in Eucalypt forest fires. This model was developed empirically from observations of hundreds of experimental fires and iteratively improving from earlier versions of forest and grassland fire models.

The Mk. 5 FFDI model takes in a limited range of inputs that are relatively easy to measure, making it a very convenient model for practical applications and one of the most widely used systems for reporting fire conditions in Australia. The model reports a Forest Fire Danger Index value based on four input variables, and calculates fire rate of spread and maximum spotting distance based on the calculated FFDI and quantity of fuel present. This system was expressed as a series of equations in 1980 by Noble et. al. [49]. The relevant equations for predicting rate of spread and spotting distance are:

$$FFDI = 2 e^{(-0.450+0.987 \ln(D)-0.0345RH+0.0338T+0.0234V)} \quad (2-11)$$

Where D is the Drought Factor, a value from 0 to 10 calculated from the quantity of rainfall and time since last rainfall. RH is relative humidity, T is temperature ($^{\circ}\text{C}$) and V is average wind speed at 10m in the open.

Using this calculated FFDI value, a rate of spread and maximum spotting distance can be calculated using equations (2-12) and (2-13) if the mass of fuel per unit area is known:

$$R = 0.0012FFDI * W \quad (2-12)$$

$$S = R(4.17 - 0.33 * W) - 0.36 \quad (2-13)$$

Where R is the rate of spread (measured in km hr^{-1}), W is the mass of fuel present in the burning area per hectare (kg ha^{-1}), and S is maximum spotting distance (km).

Although the McArthur model for rate of spread and spotting distance is widely known and the FFDI scale is popular as a means of communicating fire risk to the public, it is not a particularly accurate model of fire behaviour. Cruz and Alexander testing the performance of different wildfire rate of spread models [17], found that the McArthur Mk.5 rate of spread model had a large mean error when compared against measurements of actual fires. The Mk.5 model most often under-predicted the rate of spread, but not consistently – for example, one of the datasets tested by Cruz and Alexander (data from Burrows[50, 51]) had 66 instances of under-prediction and 34 instances of over-prediction, with no predictions falling within $\pm 5\%$ of the real observations.

Tarifa [44, 52] developed a model of firebrand transport using a semi-empirical approach. This model relies on a simplified approximation of the physics of particle motion combined with experimental data on particle characteristics. Starting with the assumption that particles behave like a point mass with the force of gravity and a drag force acting on them and constant wind conditions, the equations that describe the trajectory of a particle can be written as:

$$\alpha = \frac{1}{2} \left(\frac{\rho C_d A}{m} \right) \quad (2-14)$$

$$\frac{dw_x}{dt} + \alpha w w_x = 0 \quad (2-15)$$

$$\frac{dw_y}{dt} + \alpha w w_y - g = 0 \quad (2-16)$$

Where w is the velocity of the wind relative to the velocity of the firebrand. This would be relatively simple to solve if the particles were of constant size and mass, but the term α is not constant – the particle is burning, losing mass and changing its shape over time. To describe how α changes over time for different relative particle velocities, Tarifa performed a series of wind tunnel experiments with burning model firebrands. These experiments recorded the change in mass and dimensions over time for a burning particle in a wind field, thus giving an

empirical means of determining the value of α for a given particle at a given point in time. Some further simplifying assumptions were made with justification from the experimental results: the values of w_x and w_y were found to quickly approach their asymptotic values. This means that it is a reasonable approximation to say that the particle moves horizontally at the same speed as the wind, and that it falls vertically at its terminal velocity for its entire trajectory.

From these assumptions and experimental findings, the trajectory of a firebrand particle can be approximated if the change in terminal velocity over time from burning can be approximated, and the height at which it is shed from the convective plume is known. After an analysis of experimentally obtained terminal velocity curves, it was shown that there is a consistent relationship between the change in terminal velocity during flight and a parameter Z for the range of particles studied, shown in Figure 2-8. Since Z can be calculated based on the initial properties of the firebrand particle and time elapsed since ignition, the equations describing the trajectory of the particle can be solved if the characteristics of the particle and plume are known.




Figure 2-8: Plot of terminal velocity against parameter Z .

Although this model of firebrand transport by Tarifa is very detailed, and the experimentally obtained data is likely of good quality, it is not necessarily applicable to the problems of predicting short-range firebrand transport. The simplifying assumptions made in this model mean that it cannot account for irregularities in the wind velocity field or irregular patterns of combustion in firebrand material without modification. Irregularities in the wind field (e.g. vortices and turbulence) are likely highly important to the distribution of particles over short distances. Long and thin segments of bark that smolder or burn from one end may lose mass or change their terminal velocity at different rates to the wooden sphere and cylinder particles studied by Tarifa.

Ellis [53] produced a model of firebrand landing and spotting based on a combination of wind tunnel experiments, empirical data on the characteristics of firebrands produced from bark, and a model of wildfire plume characteristics by Raupach [54]. This model aims to combine data on the change in terminal velocity of a combusting firebrand with a model of wind velocity in the convective plume. Using this information the model can make predictions about the maximum altitude a firebrand particle may achieve and the velocity of the particle in the horizontal (downwind) axis. Bark particle terminal velocity can be calculated using an approximation found by linear regression:

$$w = 3.25a_0^{0.5} - 0.92Spin + 0.48a_0^{0.5}Species \quad (2-17)$$

Where $a_0^{0.5}$ is the square root of surface density ($\text{kg}^{0.5}/\text{m}$), $Spin$ is a value of 0 or 1 depending on whether the particle is observed to spin rapidly in the wind tunnel experiments, and $Species$ is a value determined by the species from which the particle originated. As a particle burns, the terminal velocity w changes over time due to changes in particle shape and mass. The terminal velocity of burning particles were measured in a series of wind tunnel experiments to obtain empirical data on the change in terminal velocity over time. For normal combustion behaviour, the terminal velocity at time t for a particle N can be calculated using a relationship:

$$\ln \frac{(w_t - w_{b(N)})}{(w_{t=0(N)} - w_{b(N)})} = -0.0012 \frac{t}{(V_0/S_0)\rho_0} \quad (2-18)$$

Where $w_{b(N)}$ is terminal velocity at the time a particle burns out, $w_{t=0(N)}$ is particle terminal velocity at the beginning of flight, V_0 , S_0 and ρ_0 are particle initial volume, surface area, and density. Approximations for $w_{t(0)}$ and $w_{b(N)}$ based on particle volume and surface area were found. The model then goes on to use Raupachs' equation to determine plume velocity:

$$B = \frac{g}{T_0} \frac{I_B}{\rho_a c_p} \quad (2-19)$$

Where B is the buoyancy of the source fire, g is gravitational acceleration, T_0 is 300K, ρ_a and c_p are the density and specific heat capacity of air, and I_B is the linear heat intensity of the fire (W/m). This is used to calculate vertical and horizontal plume velocity:

$$U_{yp} = 1.66B^{1/3} \quad (2-20)$$

$$U_{xp} = U_x - 0.1U_{yp} \quad (2-21)$$

Where U_{yp} is vertical plume velocity, U_{xp} is horizontal velocity within the plume, and U_x is horizontal velocity in the ambient wind field.

Using this relationship for vertical and horizontal plume velocity and the equation for particle terminal velocity as a function of time, the maximum height a particle may achieve, and the time taken for the particle to achieve that height can be calculated. The maximum horizontal distance such a particle could travel can then be calculated as:

$$X_m = (t_{hmax} U_{xp}) + (t_b - t_{hmax}) U_x \quad (2-22)$$

Where t_{hmax} is the time taken for a particle to achieve maximum altitude, and t_b is the time taken for a particle to burn out.

While Ellis' model of firebrand landing distance has a solid basis in physics, it is limited to making predictions of the maximum landing distance, rather than the frequency or density of landings over an area. Similar to other empirical models that aim to predict maximum landing distance, it is not necessarily applicable to the problem of modelling short-range firebrand distribution. However, the analysis of physics and trajectory in this model may be informative for the development of other models – in particular, investigating the maximum height

achieved across a population of particles and how that relates to the distribution of particles downwind.

2.4 Research Gap

Although there are a number of existing models that aim to predict certain aspects of firebrand behavior, each have limitations that make them less than ideal in the role of an operational model. While physics-based simulation has been shown to accurately model the distribution of firebrands in laboratory-scale experiments, it is too computationally expensive to be used as an operational model on its own. The simplified physics-based models tend to be relatively narrow in scope and lack terms to describe the distribution of firebrands under realistic wildfire conditions – where factors such as particle geometry, the size and shape of the fire, and conditions downwind of the fire can vary significantly. The distribution of firebrands in the crosswind direction is also frequently overlooked or only characterized for relatively small and narrow fires. Other models do not aim to describe a distribution at all, rather reporting only a maximum spotting distance for a given set of conditions.

Our research aims to address these gaps by developing a simple, computationally inexpensive model of short-range firebrand distribution for a range of wildfire conditions. This model is to be developed by simulating the distribution of firebrands under a range of wildfire conditions using Fire Dynamics Simulator (FDS). A key aspect of fire simulation by FDS is it computes combustion reaction between pyrolyzed fuel and oxygen. As a result, turbulence created by combustion reaction is accounted for. The simulated distribution of firebrands will then be characterized statistically, and the impact of different wildfire conditions on the shape of the distribution analyzed. This statistical analysis will form the basis of a parameterized model of firebrand landing suitable for incorporation into an operational model. The model should be able to describe the distribution of particles with a range of geometries and masses over a two-dimensional area.

Chapter 3 Model Description

Table of Contents

3.1	Introduction.....	52
3.2	Low Mach Number Approximation and Large Eddy Simulation in FDS	52
3.3	Governing Equations – Mass, Momentum, and Energy	53
3.3.1	Low Mach Number Approximation and Conservation of Energy.....	53
3.3.2	Mass Transport.....	54
3.3.3	Momentum Transport	55
3.4	Combustion Model.....	56
3.5	Lagrangian Particle Model.....	57
3.6	Solution Procedure – Time Step Selection	59
3.7	Solution Procedure – Predictor/Corrector.....	60
3.8	Summary	62

3.1 Introduction

Fire Dynamics Simulator (FDS) is a software package developed by the National Institute of Standards and Technology (NIST) for simulating fires using a physics-based approach. FDS is most often used to simulate fires within or around buildings for the purpose of evaluating the efficacy of fire safety systems design, but the physics-based modelling approach used by FDS is applicable to a wide range of fire scenarios. Importantly for the purposes of this research, the model of fluid dynamics and a model of particle transport implemented in FDS have been experimentally validated. The details of how FDS simulates certain relevant physical phenomena is described in this chapter.

3.2 Low Mach Number Approximation and Large Eddy Simulation in FDS

The foundation of physics-based modelling lies with CFD (Computational Fluid Dynamics) in which the Navier-Stokes equations are solved numerically[55]. Solving the Navier-Stokes equations for fluid motion in a complex scenario without making any simplifying assumptions – a technique known as Direct Numerical Simulation – is typically prohibitively costly from a computational perspective. In order to simulate the behaviour of fluid motion around a fire in a complex environment practically, it is necessary to make some simplifying assumptions to reduce the mathematical complexity of the problem. FDS uses a Large Eddy Simulation (LES) approach with a low Mach number approximation to achieve this reduced mathematical complexity while preserving detail in fluid flows up to a specified length scale.

The low Mach number approximation is an appropriate simplification of the Navier-Stokes equations for the purposes of fire research. In general, fire research is primarily concerned with airflows that are well under 100m/s, and usually far below that. This assumption allows for some simplification of the equations that govern conservation of energy – the particular details of which will be discussed later in this chapter.

LES is a widely-used approach for approximating the behaviour of a turbulent flow in the field of computational fluid dynamics[56]. LES simplifies the computation of fluid motion by applying a low-pass filter of width Δ to the equations governing the transport of mass, momentum, and energy. This means that the characteristics of fluid motion below a certain length scale Δ are not simulated, and instead the properties of the flow below this scale are averaged. This is a very useful simplification in practice – simulating turbulence down to a

very small scale is computationally expensive, but without necessarily having very much impact on the characteristics of fluid motion at larger scales. The properties of turbulence below the length scale Δ are approximated using a separate, simpler model referred to as the subgrid scale model (SGS).

In FDS, the LES filtering process is implemented by dividing the entire volume of the domain to be simulated into an array of rectilinear cells of known dimensions δx , δy , and δz . The filter width Δ is calculated as:

$$\Delta = (\delta x \delta y \delta z)^{\frac{1}{3}} \quad (3-1)$$

With the domain divided into fixed cells of known dimensions, the equations governing mass, momentum, and energy transport can be numerically approximated using a second-order finite difference method. Scalar quantities like temperature, pressure, and density are averaged over the whole volume of the cell. Vector quantities are averaged over the faces or edges of the cell – velocity is averaged over each face, and vorticity is averaged over each edge.

3.3 Governing Equations – Mass, Momentum, and Energy

3.3.1 Low Mach Number Approximation and Conservation of Energy

The low Mach number approximation (based on the assumption that the flow is moving at speeds well below the speed of sound) allows for some simplification of the equations that govern fluid motion, in particular equations relating to the conservation of energy. This is described in detail by Rehm and Baum[57]. This simplification is that pressure at any given point (or averaged over any given LES cell) can be described as the sum of two components: a “background” pressure $\bar{p}(z, t)$ and a “perturbation” pressure $\tilde{p}(x, y, z, t)$. Background pressure is only a function of vertical position and time, and can be approximated using the ideal gas law:

$$\bar{p} = \rho R T \sum_{\alpha} \frac{Z_{\alpha}}{W_{\alpha}} \equiv \frac{\rho R T}{W} \quad (3-2)$$

Where ρ is density, R is the gas constant, and T is temperature. Z_{α} is the mass fraction of a given gas species α and W_{α} is the molecular weight of gas species α . This “background pressure” term \bar{p} can be used in place of the actual pressure p in the state and energy equations. By making this approximation, the enthalpy of a cell h and internal energy e can be described in terms of background pressure and density: $h = e + \frac{\bar{p}}{\rho}$. The equation for conservation of

energy can then also be restated using the background pressure term \bar{p} and a new equation obtained:

$$\nabla \cdot \mathbf{u} = \frac{1}{\rho h_s} \left[\frac{D}{Dt} (\bar{p} - \rho h_s) + \dot{q}''' + \dot{q}_b''' - \nabla \cdot \dot{\mathbf{q}}'' \right] \quad (3-3)$$

Rather than solve the equation for conservation of energy explicitly, the hydrodynamics solver places a constraint on the divergence of velocity (the left hand term of the equation) such that equation (3-3) is satisfied.

3.3.2 Mass Transport

Modelling fires usually requires modelling the chemistry of combustion – oxygen reacting with a fuel to produce some products and heat. This in turn requires modelling the composition and transport of gases around fuel as it burns. To determine the quantity and composition of gases in a given grid cell at a time t , the equations that govern mass transport between cells must be solved. However, combustion is a complex process, and there may be a large number of chemical species in the air. FDS simplifies the problem of tracking different chemical species by using a “lumped species” model – each chemical species is classified as being either a “fuel” species, or a “product” species, with “air” making up the difference between them. This greatly reduces the number of transport equations that must be solved to predict the composition of gases in any given grid cell at each time step. After solving the equation for mass transport for the “lumped species”, the real composition of gases (the “primitive species”) can be calculated using a matrix of the components of each lumped species. The equation governing the transport of a lumped species is the same as that for a single species:

$$\frac{\delta}{\delta t} (\rho Z_\alpha) + \nabla \cdot (\rho Z_\alpha \mathbf{u}) = \nabla \cdot (\rho D_\alpha \nabla Z_\alpha) + \dot{m}_\alpha''' + \dot{m}_{b,\alpha}''' \quad (3-4)$$

Where Z_α and D_α are the mass fraction and diffusivity of a given species α . The terms \dot{m}_α''' and $\dot{m}_{b,\alpha}'''$ represent the production or loss of a given species from chemical sources (e.g. burning charcoal consuming O_2 and releasing CO_2). The chemical generation term with subscript b is for chemical sources from sub-grid scale particles. When applied in a Large Eddy Simulation context, the equation must be adapted to predict the composition of a grid cell at position i,j,k at time step $n+1$ based on the properties of cells at time step n . This equation becomes:

$$\frac{\delta(\rho Z_\alpha)_{ijk}^* - (\rho Z_\alpha)_{ijk}^n}{\delta t} + \nabla \cdot (\overline{\rho Z_\alpha^{FL}} \mathbf{u})_{ijk}^n = \nabla \cdot (\rho D_\alpha \nabla Z_\alpha)_{ijk}^n + (\dot{m}_\alpha''' + \dot{m}_{b,\alpha}''')_{ijk}^n \quad (3-5)$$

Where the superscript * indicates a predicted value for time step $n+1$, and superscript n represents values at the current time step. The superscript FL indicates that this value is determined with a flux limiter function – a method of interpolating values for properties of flow at cell faces. The flux limiter function is necessary to avoid numerical errors and oscillations that arise when simpler interpolation functions are used. The predicted value $(\rho Z_\alpha)_{ijk}^*$ is then modified using a “corrector” function to arrive at the final simulated value for time step $n+1$. The full solution procedure will be discussed later in this chapter.

3.3.3 Momentum Transport

The equation describing momentum in a given grid cell is an LES filtered version of the equation for conservation of momentum in a fluid. This equation can be stated as:

$$\frac{\partial \rho \bar{\mathbf{u}}}{\partial t} + \nabla \cdot (\rho \bar{\mathbf{u}} \bar{\mathbf{u}}) = -\nabla \bar{p} + \nabla \cdot \tau_{ij} + \rho g + \nabla \cdot \tau_{turb} + \mathbf{f}_b \quad (3-6)$$

Where ρ is density, $\bar{\mathbf{u}}$ is the velocity vector averaged over an LES cell, \bar{p} is pressure averaged over a cell, τ_{ij} is the viscous stress tensor, τ_{turb} is the turbulent stress tensor at sub-grid scale, g is the gravity vector, and \mathbf{f}_b is the force of drag from particles and momentum transfer from chemical processes. To solve this equation, it is necessary to find an approximation for the terms τ_{ij} and τ_{turb} . The term τ_{ij} is found using Equation (3-7):

$$\tau_{ij} = \mu(2\bar{S}_{ij} - \frac{2}{3}(\nabla \cdot \bar{\mathbf{u}})\delta_{ij}) \quad (3-7)$$

Where μ is the molecular viscosity of the fluid and δ_{ij} is the Kronecker delta function (1 when $i=j$, 0 otherwise). \bar{S}_{ij} is the strain rate tensor, defined as:

$$\bar{S}_{ij} = \frac{1}{2}(\frac{\delta \bar{u}_i}{\delta x_j} - \frac{\delta \bar{u}_j}{\delta x_i}) \quad (3-8)$$

The term τ_{turb} is more difficult to approximate. As this represents the effect of turbulence below the scale of our filter Δ , a model must be selected to calculate the effects of turbulence at scales smaller than those we aim to simulate. By default, FDS uses the Deardorff model of subgrid-scale turbulence[58]. This model treats the effects of subgrid-scale turbulence as though they were a viscous effect, relating turbulent viscosity to subgrid kinetic energy. The Deardorff model relates τ_{turb} to average flow speeds in nearby cells through a set of equations:

$$\tau_{turb} = \mu_{turb} (2\overline{S_{ij}} - \frac{2}{3} (\nabla \cdot \bar{\mathbf{u}}) \delta_{ij}) \quad (3-9)$$

$$\mu_{turb} = \rho C_v \Delta \sqrt{k_{sgs}} \quad (3-10)$$

$$k_{sgs} = \frac{1}{2} ((\bar{u} - \hat{u})^2 + (\bar{v} - \hat{v})^2 + (\bar{w} - \hat{w})^2) \quad (3-11)$$

Where μ_{turb} is the turbulent viscosity of the fluid, C_v is a constant (0.1, from literature sources [59]), and k_{sgs} is subgrid kinetic energy. The equation for k_{sgs} is the sum of the squared differences between velocity averaged over the filter width Δ (represented by values with a bar: \bar{u}) and velocity averaged over length 2Δ (represented by values with a bar and hat: \hat{u}).

3.4 Combustion Model

FDS, being a software package primarily concerned with modelling fires, includes a detailed combustion model based on a combination of chemical and physical processes. The combustion model used by default in FDS is a mixing-controlled model, where the limiting factor in the rate of combustion is the degree to which fuel and air species are mixed. This mixing-controlled combustion model uses the lumped species approach described earlier in the species transport section to model the chemical composition of each grid cell – “Fuel” species, “Air” species, and “Product” species are each lumped together to simplify the solution procedure.

At the beginning of each time step each grid cell is effectively treated as a small, turbulently mixed batch reactor that contains a combination of the three lumped species (Fuel, Air, and Products). The rate of mixing in this reactor is determined by the value τ_{mix} , which is in turn determined by the properties of the fuel, the size of the cell, and the subgrid kinetic energy in the cell. This is described mathematically as:

$$\tau_{mix} = \max(\tau_{chem}, \min(\tau_d, \tau_u, \tau_g, \tau_{flame})) \quad (3-12)$$

Each of the terms τ_{chem} , τ_d , τ_u , τ_g and τ_{flame} are used to model mixing rates constrained by different physical processes. τ_{chem} represents the lower bound of mixing time, representing the timescale of chemical transport over the thickness of the flame. τ_d represents mixing time controlled by molecular diffusion – applicable only for very small grid sizes, used in DNS. For larger length scales (such as those likely to be used in a simulation of a wildfire) τ_u and τ_g are

more likely to be applicable. τ_u is the mixing rate determined by turbulent advection, and τ_g is the mixing rate determined by buoyant acceleration. These terms are calculated as:

$$\tau_u = \frac{0.4\Delta}{\sqrt{\left(\frac{2}{3}\right)k_{sgs}}} \quad (3-13)$$

$$\tau_g = \sqrt{2\Delta/g} \quad (3-14)$$

With the value for τ_{mix} calculated using these equations, the proportion of gases in that cell that are considered “unmixed” $\zeta(t)$ is calculated as:

$$\zeta(t) = \zeta_0 e^{-\frac{t}{\tau_{mix}}} \quad (3-15)$$

FDS defaults to the assumption that $\zeta_0 = 1$, that cells are initially unmixed. Given this value for ζ_0 at time t_0 , there is sufficient information to determine the mixed fraction of each cell at later time t and thus determine the rate of combustion in each cell. The heat release rate can then be calculated based on the change in mass fraction of each species in the cell (\tilde{Y}_α) and the heats of formation of each chemical species:

$$q''' = -\rho \sum_{\alpha} \frac{(\tilde{Y}_\alpha(\delta t) - \tilde{Y}_\alpha^0)}{\delta t} \Delta h_{f,\alpha}^0 \quad (3-16)$$

3.5 Lagrangian Particle Model

Given that the primary focus of this research is on the distribution of particles, it is crucial that the model of particle motion to be used matches experimental results well and has a solid basis in physical principles. FDS uses a Lagrangian particle model to represent any sort of small particle, including static particles and liquid droplets.

The term \mathbf{f}_b in the momentum transport equation [60] represents the transfer of momentum between air and particles. This term is simply the sum of forces acting on each individual particle in a cell:

$$\mathbf{f}_b = \frac{1}{V} \sum \left[\frac{\rho}{2} C_d A_{p,c} (\mathbf{u}_p - \mathbf{u}) |\mathbf{u}_p - \mathbf{u}| + \frac{dm_p}{dt} (\mathbf{u}_p - \mathbf{u}) \right] \quad (3-17)$$

Where C_d is the drag coefficient of a particle, $A_{p,c}$ is the cross-sectional area of a particle, ρ is gas density, \mathbf{u}_p is particle velocity, \mathbf{u} is flow velocity, and m_p is particle mass. For our

purposes, the term $\frac{dm_p}{dt}$ is be taken to be zero – we are assuming that the majority of mass loss from firebrand particles occurs shortly after the particle has been shed from its source (where its temperature is highest), and that the particle has a low mass loss rate throughout the remainder of its trajectory.

The acceleration of a non-static particle is determined by finding the sum of forces acting on that particle. Here, it is simply the sum of drag forces and gravity. Using $\mathbf{F}=m_p \frac{d\mathbf{u}_p}{dt}$, the equation for particle acceleration is:

$$\frac{d\mathbf{u}_p}{dt} = \frac{\mathbf{F}}{m_p} = \mathbf{g} - \frac{1}{2} \frac{\rho C_d A_{p,c}}{m_p} (\mathbf{u}_p - \mathbf{u}) |\mathbf{u}_p - \mathbf{u}| \quad (3-18)$$

Particle motion is also impacted by sub-grid turbulence. Because sub-grid turbulence is modelled using the Deardorff model and not fully simulated, an approximation of the impact of sub-grid turbulence on particle motion can be used. This approximation adds a random fluctuation term to the gas velocity used in Equation (3-18)for each particle as it moves through each cell based on the sub-grid kinetic energy:

$$\mathbf{u} = \bar{\mathbf{u}} + u' \boldsymbol{\zeta} \quad (3-19)$$

Where $\bar{\mathbf{u}}$ is the velocity of gas flow determined by LES, u' is $\sqrt{\frac{2}{3} k_{sgs}}$ and $\boldsymbol{\zeta}$ is a vector of three normally distributed random variables with zero mean and unit variance.

In order to make practical use of these equations, we must choose a function for the drag coefficient C_d . Previous research by Wadhvani [22] investigated the performance of different drag models for modelling particle transport in FDS. Wadhvani's research found that the default FDS drag model was not ideal, and that the Haider and Levenspiel drag model [40] was the best performing model of all options tested. This drag model has been implemented in a modified version of FDS. The Haider and Levenspiel drag model calculates C_d based on the Reynolds number of the flow around the particle and particle sphericity:

$$C_d = \frac{24}{Re} (1 + A Re^B) + \frac{C}{1 + \frac{D}{Re}}, \quad Re < 2 \times 10^5 \quad (3-20)$$

The terms A, B, C, and D are functions of particle sphericity ψ :

$$A = \exp(2.3288 - 6.4581\psi + 2.4486\psi^2) \quad (3-21)$$

$$B = 0.0964 + 0.5565\psi \quad (3-22)$$

$$C = \exp(4.905 - 13.8944\psi + 18.4222\psi^2 - 10.2599\psi^3) \quad (3-23)$$

$$D = \exp(1.4681 + 12.258\psi - 20.7322\psi^2 + 15.8855\psi^3) \quad (3-24)$$

A key advantage of the Haider and Levenspiel drag model is that it is adaptable to a wide variety of geometries. By using sphericity to account for the effects of particle shape, it is possible to apply the Haider and Levenspiel drag model to any shape – although it will likely perform best for shapes that are not highly irregular.

3.6 Solution Procedure – Time Step Selection

In order to apply the LES equations for mass, momentum, and energy transfer to predict the conditions in a given grid cell at some time in the future, we need to decide how much time passes between each iteration of the solution procedure. If the time step length is too large, the results will be inaccurate due to numerical error. If the time step is too small, the simulation will become more computationally expensive for little benefit. FDS uses the Courant-Friedrichs-Lewy (CFL) condition [61] to determine time step size, as shown in equation (3-25):

$$CFL = \delta t \frac{\|\mathbf{u}\|}{\delta x} \approx 1 \quad (3-25)$$

This condition means in practice that an element of a fluid flow should not be able to travel more than one grid cell's width in a single time step. The term $\frac{\|\mathbf{u}\|}{\delta x}$ can be defined using one of three different norm functions, the default norm function used in FDS is:

$$\frac{\|\mathbf{u}\|}{\delta x} = \max\left(\frac{|u|}{\delta x}, \frac{|v|}{\delta y}, \frac{|w|}{\delta z}\right) \quad (3-26)$$

That is, we determine the largest velocity component as a fraction of grid cell dimensions and use that component to determine the CFL criterion.

An additional constraint on time step size is determined by conservation of mass. A fluid element with mass ρV does not gain or lose mass over time. It can be shown that for a fluid

element of constant mass, the divergence of velocity over that element is related to the volume of that element by:

$$\nabla \cdot \mathbf{u} = \frac{1}{V} \frac{dV}{dt} \quad (3-27)$$

The divergence of velocity over the volume of the fluid element determines whether the element is being compressed or expanding. For a time step of size Δt , we can say that volume at time step $n+1$ is related to volume at time n by:

$$V^{n+1} = V^n + \Delta t V^n (\nabla \cdot \mathbf{u})^n \quad (3-28)$$

Since negative volumes are impossible, we can define a new constraint on the length of the time step that ensures volume at time step $n+1$ is positive. This constraint is:

$$\Delta t < -(\nabla \cdot \mathbf{u})^{-1} \quad (3-29)$$

We can incorporate this constraint into a modified version of the CFL condition, producing the equation:

$$CFL = \delta t \left(\frac{\|\mathbf{u}\|}{\delta x} + |\nabla \cdot \mathbf{u}| \right) \quad (3-30)$$

FDS aims to keep this modified CFL condition value in a range between 0.8 and 1. At each time step, the CFL condition is checked. If the time step is too large and the CFL value is greater than 1, the calculations are repeated with a smaller time step equivalent to a CFL condition value of 0.9. If the CFL value is less than 0.8, the time step is increased by 10%. This process ensures that the time step is close to an optimal value throughout the entire computation.

3.7 Solution Procedure – Predictor/Corrector

Having described the relevant equations for mass, momentum, and energy, we can describe the procedure FDS uses to predict conditions at time step $n+1$ from conditions at time step n . This procedure uses a predictor-corrector approach, where a predictor value (marked by a superscript ‘*’) is calculated based on conditions at time step n , and that predictor value is adjusted using a corrector function to produce the final value for physical conditions at time $n+1$.

The process begins with a known set of information about a particular grid cell at time step n . This information includes the density ρ^n , the lumped species mass fractions Z_α^n , velocity vector \mathbf{u}^n , and Bernoulli integral H^n . For every cell in the domain, a background pressure \bar{p}^n is known. Temperature is found using the equation of state. This information is used to produce the predictor values, which are then used to produce the final corrected values. The solution procedure is described below:

Predictor:

1. Estimate ρ , Z_α and \bar{p} at the next time step using an explicit Euler step. The species mass density is estimated as:

$$\frac{(\rho Z)_\alpha^* - \rho^n Z_\alpha^n}{\delta t} + \nabla \cdot \rho^n Z_\alpha^n \mathbf{u}^n = \nabla \cdot (\rho^n D_\alpha^n \nabla Z_\alpha^n) + (\dot{m}_\alpha''' + \dot{m}_{b,\alpha}''')^n \quad (3-31)$$

Note that the value of $\nabla \cdot \rho^n Z_\alpha^n \mathbf{u}^n$ in this equation may be modified by the use of a flux limiter function to minimize errors in the final result.

2. Compute the density from $\rho^* = \sum_\alpha (\rho Z)_\alpha^*$ and mass fractions from $Z_\alpha^* = (\rho Z)_\alpha^* / \rho^*$.
3. Compute the temperature T^* from the equation of state.
4. Compute the divergence predictor $(\nabla \cdot \mathbf{u})^*$ using estimated quantities from the previous steps:

$$\nabla \cdot \mathbf{u} = \frac{1}{\rho h_s} \left[\frac{D}{Dt} (\bar{p} - \rho h_s) + \dot{q}''' + q_b''' - \nabla \cdot \dot{\mathbf{q}}'' \right] \quad (3-32)$$

5. Solve the Poisson equation for the pressure term:

$$\nabla^2 H^n = - \frac{(\nabla \cdot \mathbf{u})^* - \nabla \cdot \mathbf{u}^n}{\delta t} - \nabla \cdot \mathbf{F}^n \quad (3-33)$$

6. Estimate the velocity predictor \mathbf{u}^* using the results obtained from the previous step:

$$\frac{\mathbf{u}^* - \mathbf{u}^n}{\delta t} + \mathbf{F}^n + \nabla H^n = 0 \quad (3-34)$$

7. Check that the time step length δt satisfies the modified CFL condition. If the time step is too large, repeat the procedure with a smaller time step length so that the condition is met. Otherwise, proceed to the corrector step.

Corrector:

1. Correct the transported species mass densities using the equation below:

$$\begin{aligned} \frac{(\rho Z)_\alpha^{n+1} - \frac{1}{2}(\rho^n Z_\alpha^n + \rho^* Z_\alpha^*)}{\delta t/2} + \nabla \cdot \rho^* Z_\alpha^* \mathbf{u}^* \\ = \nabla \cdot (\rho^* D_\alpha^* \nabla Z_\alpha^*) + (\dot{m}_\alpha''' + \dot{m}_{b,\alpha}'')^n \end{aligned} \quad (3-35)$$

2. Compute the density $\rho^{n+1} = \sum_\alpha (\rho Z_\alpha)^{n+1}$ and mass fractions from $Z_\alpha^{n+1} = (\rho Z)_\alpha^{n+1} / \rho^{n+1}$.
3. Compute the temperature T^{n+1} from the equation of state.
4. After the corrector step for mass and species transport, source terms are computed and stored. This includes terms for heat release rate per unit volume (\dot{q}'''), the net absorption of thermal radiation ($\nabla \cdot \dot{\mathbf{q}}''$), and source terms for mass (\dot{m}_α'''). These terms are stored and will be applied in the predictor and corrector terms of the next time step.
5. Compute the divergence $(\nabla \cdot \mathbf{u})^{n+1}$ using the corrected results obtained from the previous steps.
6. Compute the pressure using the estimated quantities:

$$\nabla^2 H^* = - \frac{(\nabla \cdot \mathbf{u})^{n+1} - \frac{1}{2}(\nabla \cdot \mathbf{u}^n + \nabla \cdot \mathbf{u}^*)}{\delta t/2} - \nabla \cdot \mathbf{F}^* \quad (3-36)$$

7. Correct the velocity at the next time step:

$$\frac{\mathbf{u}^{n+1} - \frac{1}{2}(\mathbf{u}^n + \mathbf{u}^*)}{\delta t/2} + \mathbf{F}^* + \nabla H^* = 0 \quad (3-37)$$

3.8 Summary

This chapter has given an overview of the mathematical and physical basis of the core functionality of FDS where it is relevant to the simulations we aim to run. The mathematical basis of the model is shown to be sound, and the low Mach number assumption is reasonable for the purposes of wildfire modelling. The Lagrangian particle model implemented in FDS has been modified with an alternative method for calculating drag coefficients (the Haider and Levenspiel model), but the underlying physics and mathematics of how particle motion is handled in FDS has not been modified. Some components of FDS, such as those relating to pyrolysis, are not described in detail here because they are not particularly relevant to the research being undertaken.

Chapter 4 Scenario Development and Sensitivity Analysis

4.1 Selection of scenarios

In order to obtain useful results from a physics-based simulation of firebrand transport, the scenarios we choose to simulate must be well-designed and properly specified. This means that we must design our scenarios to eliminate or minimize the impact of errors caused by the limitations of grid size, domain size, boundary conditions, and other computational limitations. We must also be careful to choose appropriate input values for parameters relating to fire behaviour, firebrand properties, and atmospheric conditions.

Due to the high computational cost of CFD simulation, we must choose a limited selection of parameters to vary between scenarios. Each additional parameter we choose to study increases the computational cost of these scenarios multiplicatively, meaning it is only practical to study the most impactful parameters. These parameters are:

- Wind Speed (Four wind speeds selected to represent moderate to extreme fire conditions)
- Level of forest cover over downwind terrain (Forested and Unforested)
- Aspect ratio of flame width to depth (Three aspect ratios chosen to represent wide to narrow fires – 4:1, 2:1, 1:1)
- Heat Release Rate (HRR) per Unit Area (Two values chosen – equivalent to 100% fuel consumption and 66% fuel consumption)

This choice of parameters is informed by existing operational fire models where wind speed, forest coverage, and fuel load (and therefore heat release rate) are common inputs to models [25]. Flame width to depth aspect ratio is a unique parameter, where real fires can have highly variable widths depending on unique circumstances in local geography and wind conditions. In our case, we choose to simulate a range of widths due to how these widths are expected to alter the characteristics of the convective plume[62].

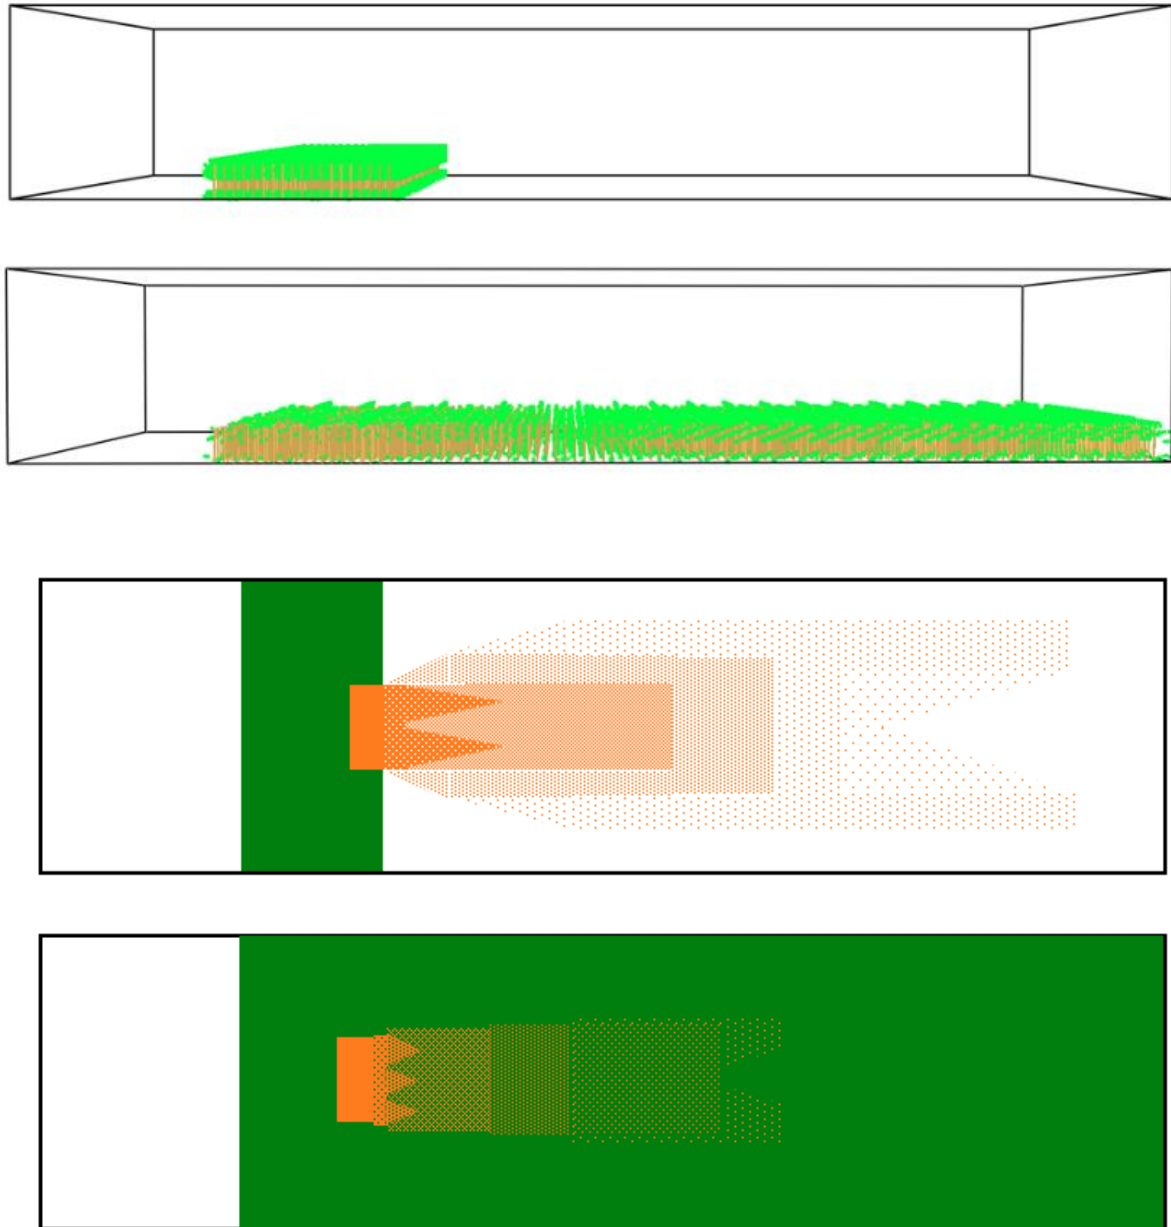


Figure 4-1: Diagram of Forested and Non-forested domains, and an example of firebrand landing density in each case (not to scale).

By specifying the characteristics of vegetation in our scenario, we are indirectly specifying the mass loss rate (MLR) and heat of combustion of fuel present for the fire to burn. This then determines the conditions of the simulated combustion reaction between pyrolyzed fuel (specified by MLR) and oxygen. As a result of this exothermic reaction, heat is released and a convective plume including turbulent flame is produced. The quantity of firebrand production is linked to HRR and firebrands are simultaneously lofted by the plume and driven by the winds. Given how these properties are interlinked, we must begin by specifying some initial parameters for the vegetation in our scenario.

The McArthur Mk.5 Forest Fire Danger Index (FFDI)[63] is a commonly used system for rating the risk of a bushfire given certain parameters: temperature, relative humidity, wind speed, and drought factor (reflecting dry fuel availability). While the Mk.5 FFDI is not the most accurate system for predicting fire rate of spread, it provides a simple scale for informing the public about the risk of bushfire on a given day with very straightforward calculations. The Mk.5 FFDI is calculated by:

$$FFDI = 2e^{-0.45+0.987\ln(DF)-0.0345RH+0.0338T+0.0234U} \quad (4-1)$$

Where DF is Drought Factor (scale of 1 to 10), RH is Relative Humidity, T is temperature in Celsius, and U is wind speed at 10 m above the ground in km/hr. This provides a useful starting point for choosing atmospheric conditions that are likely to reflect severe to extreme fire conditions. By selecting ambient temperatures, humidity, and wind speeds that generate a range of high FFDI values, we can begin to narrow the range of parameters for scenarios. For our purposes, four FFDI values were chosen: 80, 100, 120, and 150. By specifying a fixed temperature, drought factor, and relative humidity, these FFDI values correspond to specific wind speeds as described in Table 1.

Table 1: Parameters for FFDI values

FFDI	Temperature	RH	Wind Speed
80	40 °C	10%	37km/hr
100	40 °C	10%	46km/hr
120	40 °C	10%	54km/hr
150	40 °C	10%	64km/hr

With these critical atmospheric conditions defined, we can go on to use more accurate models of vegetation growth and bushfire behaviour to calculate parameters relating to vegetation.

Project Vesta [36] was a CSIRO investigation into the fuel structure and fire dynamics of bushfires in Australian Eucalypt forest. This project provides several useful correlations for predicting fuel conditions in eucalypt forests based on age. These correlations are based on experimental observations of a large number of individual plots of eucalypt forest, with classification of fuel load in five layers: surface, near-surface, elevated, intermediate, and

overstorey. Fuel loads in the surface and near-surface layer were found to accumulate in a predictable pattern over time, described by a regression model as:

$$\text{Near Surface Fuel Load} = ss(1 - e^{kt}) \quad (4-2)$$

where ss and k are regression parameters, and t is the age (in year) of the fuel bed. This model provides a basis for predicting the maximum near-surface fuel load that may accumulate in a reasonable period of time in a real eucalypt forest, with approximately 18 tons/ha as an upper limit for a forest that has not experienced a fire in the last ~20 years.

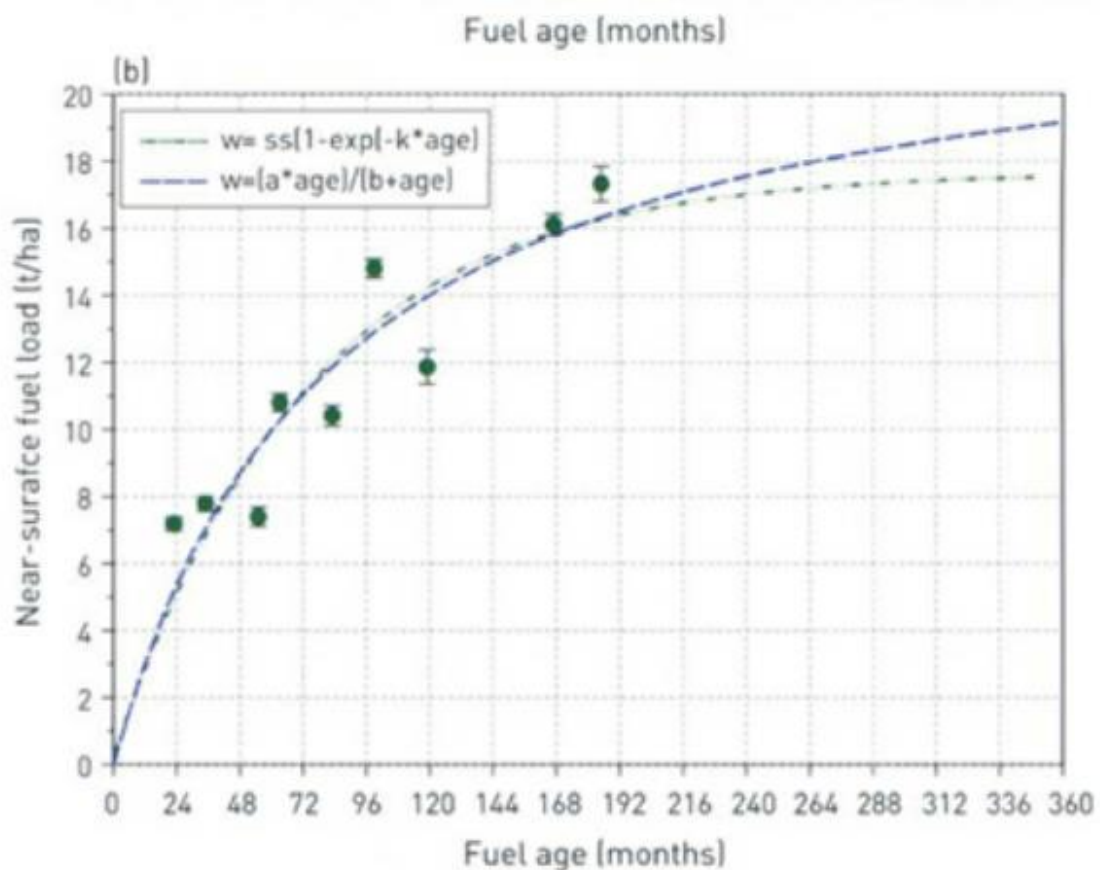


Figure 4-2: Near-surface fuel load vs. age, from Project Vesta [36]

With the mass of near-surface fuel per unit area in our scenario now specified, choosing an appropriate heat of combustion and flame residence time will yield a maximum potential HRR Per Unit Area (HRRPUA). Cone calorimetry experiments performed by Wadhwani at VU [64] provide a value of 15.72 MJ/kg for surface fuels in Eucalypt forests. With a residence time of ~30s (similar to that observed in Figure 4-3) and assuming full combustion of the near-surface fuel load, we can expect a HRRPUA of 943 kW/m².

However, as observed in Project Vesta [36], fuels in the elevated, intermediate and overstorey layers also contribute to the total mass of fuel burned in a bushfire. Fire rate of spread was observed to be positively correlated with the elevated and intermediate fuel hazard scores (a score based on researcher observation of a variety of elevated fuel characteristics). Although a simple regression model is not available to estimate the fuel load in these layers, a table of correlations between these scores and slope-adjusted flame rate of spread is available. To account for the contribution of fuels in these layers to the overall rate of spread, an additional mass of fuel to the 18 tons/ha of near-surface fuels was added such that peak flame temperature in the scenario matched that of Project Vesta observations. This additional mass from the elevated and intermediate layers contributes approximately 30% of total fuel burned in our scenario, increasing HRRPUA to 1362 kW/m² for 100% fuel consumption and 899 kW/m² for 66% fuel consumption.

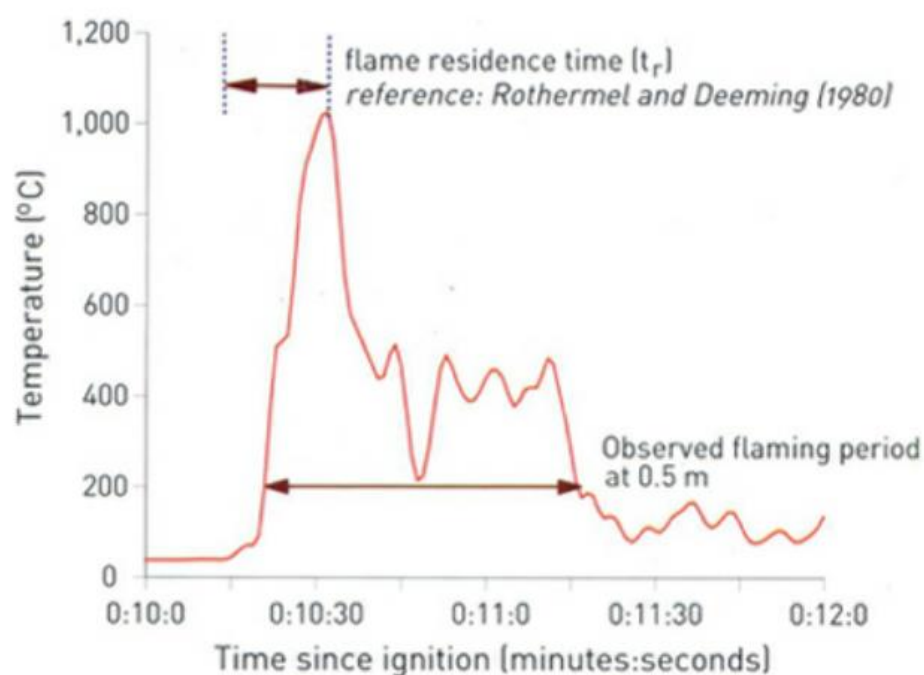


Figure 4-3: Flame temperature vs. time observed in Project Vesta controlled burning experiments [36]

In addition to defining the mass of fuel to be burned, Project Vesta provides several correlations for fuel hazard scores based on fuel age. These hazard scores are used as an input for the Project Vesta fire Rate of Spread (ROS) model alongside atmospheric conditions to predict a steady-state ROS of a fire. The Vesta ROS model consists of three equations that together account for fuel and atmospheric conditions, moisture content, and the effect of slope:

$$R_{ss} = R_t + b_0(U_{10} - U_t)^{b_1} \cdot e^{(b_2 S_{fhs} + b_3 NS_{fhs} + b_4 NS_h)} \cdot \phi M_f \cdot \phi S_f \quad (4-3)$$

$$\phi M_f = (M_f^{-1.495})/0.0545 \quad (4-4)$$

$$\phi S_f = e^{0.069\theta} \quad (4-5)$$

Where:

- R_{ss} is the steady-state rate of spread
- R_t is a threshold rate of spread (ROS at 5km/hr wind speed)
- U_{10} is wind speed 10m above the canopy
- U_t is a threshold wind speed (5km/hr)
- S_{fhs} is surface fuel hazard score
- NS_{fhs} is near-surface fuel hazard score
- NS_h is near-surface fuel height
- b_0 to b_4 are coefficients for wind speed and fuel hazard scores.
- ϕM_f and ϕS_f are functions that correct for fuel moisture (M_f) and slope (θ)

Fuel moisture content can be calculated from relative humidity using relations developed by Matthews et al. [65]. These fuel moisture calculations give a fuel moisture of ~3.5% for the atmospheric conditions we have specified. We can also assume that our fire occurs on flat ground, with $\theta = 0$, so we can set the slope modifier function to 1. Knowing these values, we can then use the correlations for fuel hazard score and near-surface fuel height from Project Vesta for a forest of 20 years age to calculate steady-state ROS for each of our scenarios. These ROS values, combined with our fixed flame residence time, also produce a fixed flame depth for each scenario.

Table 2: Various ROS, fireline depth and fire residence time for different velocities

Wind Speed	ROS	Fireline depth	Residence time
37km/hr	1.08 m/s	32.4 m	30s
46km/hr	1.38 m/s	41.3 m	30s
54km/hr	1.58 m/s	47.3 m	30s
64km/hr	1.86 m/s	55.8 m	30s

4.2 Thermo-physical properties of vegetation and firebrand characteristics

The proper specification of vegetation and firebrand parameters is critical for this research. Due to the importance of these factors, an experimental basis for the size, shape, density and quantity of firebrands was needed. This area of firebrand generation and characterization has been investigated using vegetation burning experiments, such as those conducted by Manzello et al [30, 66], and Hudson et al [33], investigating common American species such as Douglas Fir, Grand Fir, Loblolly Pine, Ponderosa Pine, and Western Juniper. However, these experimental studies do not provide the specific data we require for Australian environments – in particular, data relating to the physical characteristics and total quantity of firebrands released per kilogram of mass lost for Eucalyptus fuels. Recently, Wickramasinghe et al [31] has proposed a method by which results for these American species can be used to calculate parameters for Eucalypt forests.

In the proposed method of Wickramasinghe et al, a series of correlations are used for firebrand generation based on wind speed, moisture content, fuel load, species of fuel present, and fuel heat of combustion to produce a value for the total number of firebrands produced per megawatt of heat output. Given that the sub-canopy wind speed in these scenarios is low in all cases (around 2m/s at the highest wind speeds), and the bulk of the fuel is close to the ground, there was little variation in firebrand generation rate per megawatt across different wind speeds. Values for firebrand generation per MW fell between 6.1 pcs/MW/s and 7.2 pcs/MW/s.

Having defined the quantity of firebrands produced per megawatt in each case, the physical dimensions and densities of the firebrands must be defined. Some literature sources for certain characteristics of Eucalypt firebrands are available (such as histograms of particle masses)[35] but lack other key criteria, such as detailed descriptions of particle geometry or images of large sets of particles. As such, the physical dimensions of firebrands in our simulations were based on digital images of firebrands obtained from vegetation burning experiments by El Houssami et al. [67] These particles include a variety of bark, twig, and leaf fragments of different geometries and sizes, giving a range of particle species that will exhibit different transport characteristics.




Figure 4-4: Image of firebrands obtained by Houssami et al. [67]

These particles exhibit significant variation in size and shape, but for the purposes of transport modelling in FDS with the Haider & Levenspiel drag model, each one must be categorized as belonging to one of three particle geometries: cylindrical, spherical, or cartesian (rectangular). Using results obtained previously by Wickramasinghe et al [31], these particles were divided into 42 particle categories, each belonging to one of these three geometries. The range of particle sizes for each of these geometries is summarized in Table 3.

Table 3: Range of particle dimensions for different geometries, *diameter, **radius

Geometry	Width	Length	Depth
Cylindrical	1.25mm-14.12mm*	8.77mm-91.30mm	-
Spherical	2.89mm-15.66mm**	-	-
Cartesian	5.43mm-23.26mm	8.08mm-36.71mm	3mm

We can then compare the dimensions of the firebrands in Table 3 with data relating to Eucalypt firebrands from literature. Literature sources are available for the projected area of firebrands from Eucalypt species[35]. As shown in Figure 4-5, real Eucalypt firebrands are predominantly composed of particles in the range $10\text{-}50\text{mm}^2$ (65-84% of particles) with the remainder split between the $50\text{-}100\text{mm}^2$ and $100\text{-}500\text{mm}^2$ ranges, and a few outliers beyond that. This agrees well with the dimensions of firebrands obtained by Wickramasinghe et al, indicating that this range of particle dimensions is likely suitable for our purposes.




Figure 4-5: Histogram of firebrand projected areas for various species – Eucalypt results are EUC01 and EUC02. From [35].

With these physical parameters for atmospheric conditions, vegetation, fire, and firebrands set, preliminary simulations of fires can be done to obtain useful results that will assist in properly specifying the domain.

It should be noted that the FDS model of Lagrangian particles does not allow for particles to collide with each other. This means that the interaction of firebrand particles with other small particles of vegetation (e.g. leaves) present in the burning region is limited to interactions mediated through airflow (such as drag, and thus resultant changes in the direction of the wind, plume, and vortices). As such, the vegetation morphology in the domain is quite simple, and the interaction of firebrand particles with this vegetation is limited.

4.3 Domain size and boundary conditions

Choosing an appropriate computational domain with the correct boundary conditions is essential for the development of our scenario in FDS. The domain must be sized and specified such that all the necessary transport and atmospheric phenomena relevant to short-range firebrand transport can be adequately simulated. In addition to this, the domain must be kept small enough so that the computational cost of running these simulations is not excessive. The boundary conditions of our scenario must be chosen so that the limited size of the domain does not influence the result significantly.

4.3.1 Domain Size

There are three major requirements that must be taken into consideration when choosing an appropriate domain size. First, the domain must be large enough to capture all the relevant phenomena that we aim to study. Second, the amount of computing power required to run each instance of the simulation must be kept within reasonable limits. Third, domain height should be large enough to allow firebrands lofted in a high arc to stay within the domain for the entire length of their trajectory. Issues of computational power are discussed in the grid convergence section of chapter 4.

In order to capture all relevant phenomena for the process of firebrand transport, we must first have a properly developed atmospheric boundary layer in which our firebrands may be transported. This in turn requires that we size our domain so that there is enough space immediately after the inlet for such conditions to develop. To ensure that sufficient distance after the inlet is given, a series of time-averaged wind profiles were computed to ensure that the profile was fully developed before the edge of the forested area. After letting the wind profile develop for 450 seconds, wind profiles were taken at 10m intervals between 10m and 180m from the inlet and averaged over 150 seconds. The position of each interval along the ground is indicated in Figure 4-6:




Figure 4-6: Measurement intervals before forest canopy

It was found that between each of these time-averaged profiles, average wind speed varied by less than 1%, with an average difference of 0.761% between intervals. However, there was a clear trend towards slightly lower speeds at low altitudes as the wind profile approached the

edge of the forest. Using results from intervals further from the forest (thus closer to the inlet) reveals that the wind profile appears to be well-developed within the first 50m of the inlet, with negligible changes between each interval (an average of 0.32% difference between intervals). Taking this result into account, it is clear that 200m would be more than sufficient distance from the inlet for the wind profile to have become fully developed.




Figure 4-7: Convergence of average wind velocity profiles in range 10-50m. Circled point is 10.3m/s at 10m, where U_{10} was specified.

After the wind field is developed, an area must be defined for foliage, fire, and particle insertion. This area must be large enough that a realistic sub-canopy wind field can be developed, as well as contain the entire space our fire might occupy. Based on preliminary calculations of fire rate of spread, the space required for fires will be at least 100m deep and approximately 200m wide. An additional 100m of foliage depth was included for the development of a realistic sub-canopy wind field, and an additional 150m width was given on each side to allow for the possibility of firebrand transport with a wide lateral dispersion. Thus, an area of 200m depth and 500m width of the domain was populated with foliage in the open-field case. This area of foliage was extended throughout the entire depth of the domain after the initial 200m for the forested case.




Figure 4-8: Illustration of domain with foliage and dimensions.

After the area to be burned has been defined, the domain must be extended so that all firebrand transport phenomena can be captured. However, the domain must also be kept small enough so that it does not become impractically computationally expensive to run. Each of these requirements places certain constraints in terms of what dimensions we can set for our domain.

For our purposes, “short range” firebrand transport is defined to occur within 750m of the fire front – this choice of 750m is in accordance with the definition by Cruz et. al given in the CSIRO/AFAC Guide to Fire Spread Models[25]. This means that at minimum, we must extend the domain an additional 750m from the edge of the fire. In addition to this, an extra 50m was added to the end of the domain, for a total of 800m of domain length downwind of the fire front. This domain length ensures that all short-range firebrand transport can be recorded, and the landing density of firebrands beyond that distance can be confirmed to be extremely low. Extending the domain beyond this length would increase computational costs for very little benefit in terms of observed firebrand transport behaviour, as the plot of firebrand distribution presented in Figure 4-9: plot of firebrand landing density in 54km/hr wind conditions, 4:1 aspect ratio. Edge of fire line is at 400m. Figure 4-9 indicates extremely low landing densities at those distances. Moreover, short-range firebrands traveling further than 800m from their point of origin are not likely to have sufficient energy to start a new fire.

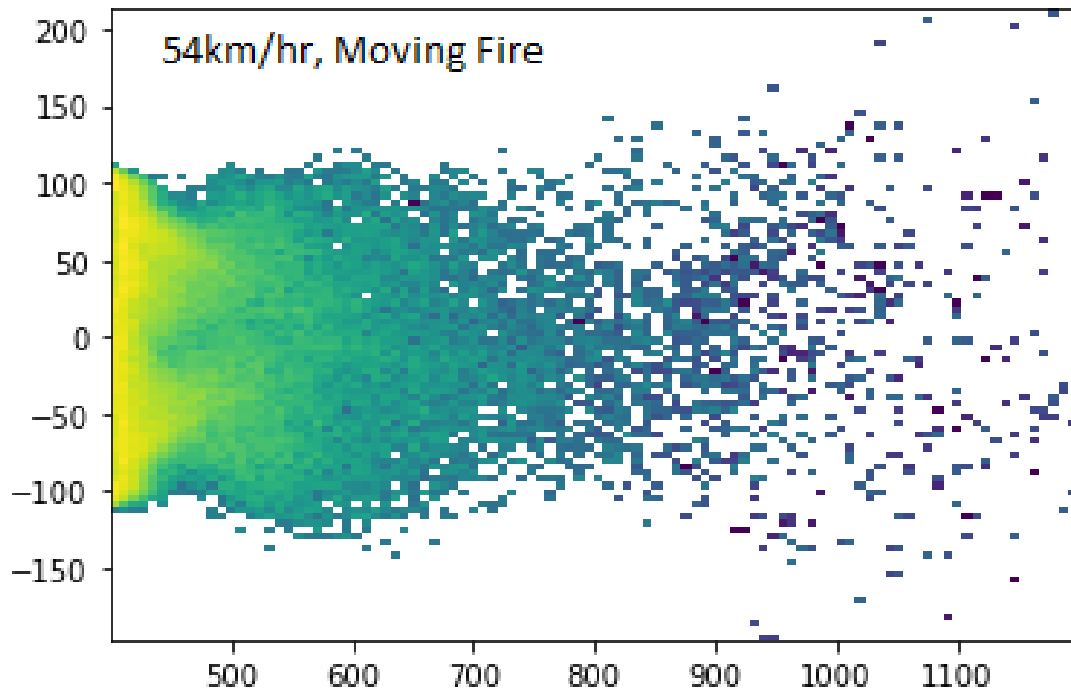


Figure 4-9: plot of firebrand landing density in 54km/hr wind conditions, 4:1 aspect ratio. Edge of fire line is at 400m.

The width of the domain is maintained at 500m throughout its entire length. The reasoning for this width was discussed previously, as the width required for the largest fire with sufficient room for a particularly wide particle distribution.

The height of the domain was chosen so that firebrands lofted in a high arc that could land within 750m do not escape the domain. An initial estimate of 200m domain height was selected, and several fires were simulated in a domain with this height. Based on these preliminary simulations, less than 0.043% of firebrands ever reached a height of greater than 195m. While some particles may escape the domain, the computational cost of extending the domain beyond 200m height to capture the behaviour of this extremely small population of particles was not considered worthwhile. A sensitivity analysis of domain height was conducted to confirm that velocity profiles at lower altitudes were not significantly affected by the choice of domain height, the results of which are shown in Figure 4-10:




Figure 4-10: Change in wind velocity profiles at 100m from the inlet across 175m, 200m, and 225m domain heights.

Note that in the 175m domain height case, there is a reduction in wind speed close to the top of the domain. This is due to errors caused by the open boundary condition and flows moving parallel to the top of the domain – discussed in section 4.3.3. Aside from this effect, the velocity profiles are extremely close in each case.

4.3.2 Inlet Conditions

The atmospheric conditions throughout the domain depend on a variety of properties chosen at the inlet. There are several methods of specifying these inlet conditions, using different models of wind velocity gradients and temperature conditions in the atmospheric boundary layer. Of these various methods, two were considered for use in our scenario – the wall of wind method with synthetic eddies[68], and Monin-Obukhov Similarity Theory (MOST) [69].

The wall of wind method is the simpler of the two methods of generating a profile of wind velocity at the inlet of our domain. In this method, the atmosphere in the domain is initialized with zero initial velocity and a constant temperature. The inlet of the domain then allows a “wall of wind” to travel through the domain, with a specified power law scaling the inlet velocity with altitude. Given this method starts with zero wind velocity throughout the domain,

a realistic wind profile is not developed until enough time has passed within the simulation for the wind to have completely swept through the domain several times over.

In addition to this, the wall of wind method does not have any turbulent eddies present in the initial wind field, meaning they must be specified manually. This is achieved via the synthetic eddy method, where eddies of a known size and average velocity perturbation are inserted over the area of the inlet surface. These eddies are specified using equations (4-6 and (4-7:

$$L_{eddy} = 3\Delta x \quad (4-6)$$

$$N_{eddy} = \frac{L_y L_z}{L_{eddy}^2} \quad (4-7)$$

Where L_{eddy} is the length of each eddy inserted at the inlet, N_{eddy} is the number of eddies to be inserted, Δx is the grid size in the x direction, and L_y and L_z are the dimensions of the inlet in the y and z directions.

An alternative method to the wall of wind with synthetic eddies is to use Monin-Obukhov Similarity Theory. By specifying a value for the Obukhov length (L) and surface roughness length (z_0), temperature and wind velocity profiles can be determined for the initial conditions present in the domain. This method specifies atmospheric conditions throughout the entire domain at the instant the simulation begins, not just at the inlet.

However, as the Monin-Obukhov method specifies velocity throughout the whole domain, certain issues arose which made it inadvisable for our purposes. Wind velocity in heavily obstructed areas – such as within the forested areas of the domain – was considerably too high, unlike values suggested in literature or those produced using the wall of wind method. This leaves only the wall of wind method as an appropriate option for generating wind at the inlet of the domain. There is also some discussion in literature [70] on issues that arise from the use of MOST in LES with fine grid sizes and large roughness elements, which would advise against its use in our scenarios.

4.3.3 Boundary Conditions

The boundary conditions specified at the exterior surfaces of the domain must be chosen so that they do not impact the phenomena we wish to investigate. FDS has three options for external domain boundaries aside from solid walls [69] – open, mirror, and periodic. Open boundaries allow gas and particle transport through the boundary, with some assumptions made about conditions beyond the boundary (such as pressure and temperature terms outside the

domain). Mirrored boundaries act as though the conditions on the exterior side of the boundary are identical to those on the interior side – meaning there is no flux through a mirrored domain boundary. Periodic boundaries copy their conditions from another periodic boundary on the other side of the domain – effectively allowing a user to specify a domain that repeats itself after a certain distance.

Ideally for our purposes, all exterior boundaries aside from the ground would be “open” – allowing wind and particles to flow freely through them. However, the open boundary condition in FDS has certain undesirable behaviours when wind flows parallel to the boundary, rather than through it. This is best understood by considering the two different expressions for H at an open boundary used for outgoing and incoming flows:

$$H_{\frac{1}{2},jk} = \begin{cases} \frac{\tilde{p}_{ext}}{\rho_{ijk}} + \frac{1}{2}(\bar{u}_{1,jk}^2 + \bar{v}_{1,jk}^2 + \bar{w}_{1,jk}^2) & \text{outgoing} \\ \frac{\tilde{p}_{ext}}{\rho_{\infty}} + \frac{1}{2}(u_{\infty}^2 + v_{\infty}^2 + w_{\infty}^2) & \text{incoming} \end{cases} \quad 4-8$$

Incoming flows calculate H using values for ρ , u , v , and w at infinity, and outgoing flows use values from the cell adjacent to the boundary. This means that small changes in the direction of the wind flowing parallel to an open boundary can cause the pressure solver to switch between using values at infinity and values from an adjacent cell, which can be significantly different.

Due to these pressure calculation issues for wind travelling parallel to an open boundary, it was observed that in practice using the “open” condition for boundaries on the side of the domain tended to create a drift in particle motion towards the positive y direction. These issues were successfully resolved by using a “mirror” boundary condition instead, where pressure on either side of the boundary is assumed to be equal, as shown in Figure 4-11 below.




Figure 4-11: Comparison of firebrand distribution histograms in 46km/hr wind conditions. (a) Distribution with open side boundaries and (b) Distribution with mirror boundaries

Despite the majority of flow travelling parallel to the top surface of the domain, this boundary cannot be set to “mirror” as the convective plume needs to travel through it to model realistic transport behaviour. This boundary surface is set to “open”, and the pressure fluctuations caused by this boundary condition have been observed to cause a gradual slowing of wind at

the very top of the domain in comparison to mirrored test cases. This effect is minimal beyond the final few meters of the domain in the vertical axis.

The boundary conditions at the entrance and exit of the domain are set to open. In this case, the open boundary condition performs as expected, as flows are directly perpendicular to the surface of the boundary.

4.4 Grid Convergence Analysis

After specifying the dimensions of the domain to be simulated, it is necessary to define the grid size to be used for CFD Simulation. FDS allows the user to specify grid size throughout the entire domain, or split the domain into many component meshes, each of which can be assigned their own grid size (though this size must be a multiple of two of an adjacent grid size).

The selection of an appropriate grid size is essential for the accuracy and computational feasibility of CFD analysis. A process of gradual refinement, known as a grid convergence analysis, allows us to examine the impact of changes in grid size on the outcome of our simulations. As grid size is reduced, the results of each simulation will tend to converge, as errors caused by coarser approximations to the Navier-Stokes equations are eliminated. Further reductions in grid size after convergence has been obtained yield negligible improvements in accuracy. To be meaningful, any results obtained from a CFD fire model should be grid converged.

It is therefore necessary to quantify the impact of changes in grid size on the wind field to ensure that our results are grid converged. Finer grid size comes with a rapidly scaling computational cost to load and process each cell in the domain, as well as the requirement of proportionally smaller time steps to meet the CFL condition:

$$CFL = \delta t \frac{\|\mathbf{u}\|}{\Delta} < 1 \quad (4-9)$$

Where δt is the size of our time step, and $\frac{\|\mathbf{u}\|}{\Delta}$ is the norm of the velocity vector:

$$\frac{\|\mathbf{u}\|}{\Delta} = \frac{|u|}{\delta x} + \frac{|v|}{\delta y} + \frac{|w|}{\delta z} + |\nabla \cdot \mathbf{u}| \quad (4-10)$$

By performing a grid convergence analysis, we can choose a grid size that is both practical to run and that ensures convergence. Starting with a coarse grid, a simulation is run with some key quantity chosen as an output. The simulation is then re-run with a finer grid, and the difference in the key quantity between the coarse and fine cases is used to calculate a Grid Convergence Index (GCI)[71], calculated as:

$$GCI = \frac{F_s |\varepsilon|}{(r^p - 1)} \quad (4-11)$$

Where F_s is a safety factor (a multiplier of the relative error term, chosen so that our measured GCI is a conservative estimate of actual grid convergence), ε is the relative error of the chosen quantity, r is the ratio of grid refinement, and p is the order of convergence of the simulation[72]. This formula accounts for the difference in grid size between each case, as well as the order of the numerical scheme used to approximate the solution to the Navier-Stokes equations. FDS uses an explicit second-order numerical method to calculate each time step, so the order of convergence p is two. The grid refinement ratio is specific to each pair of cases being compared, and is the ratio of grid dimensions between the larger and smaller grids (e.g. if comparing a 2m and 1.5m grid, $r = 2/1.5$). F_s is generally chosen based on the circumstances of the grid convergence study – for our purposes, $F_s = 3$ is a suitable (if conservative) value supported by literature examples[72]. The calculation of the relative error term depends on the specific quantity being measured for the grid convergence study. Lower the GCI, greater the grid convergence. A low GCI is indicative of good grid convergence.

To perform a grid convergence analysis of our scenarios, a series of simulations with 2m, 1.5m, and 1m grid sizes containing only empty space and foliage were run, allowing wind to flow through the domain and establish a well-developed velocity profile. Two locations were chosen to measure grid convergence: 150m from the inlet before the forest canopy, and 300m from the inlet, above the centre of the forest canopy.

For measurements at 150m, wind velocities were recorded at the centre of every grid cell along a vertical column at a fixed location 150m downwind from the inlet surface. These velocities were averaged over a period of 120s to produce wind profiles such as shown in

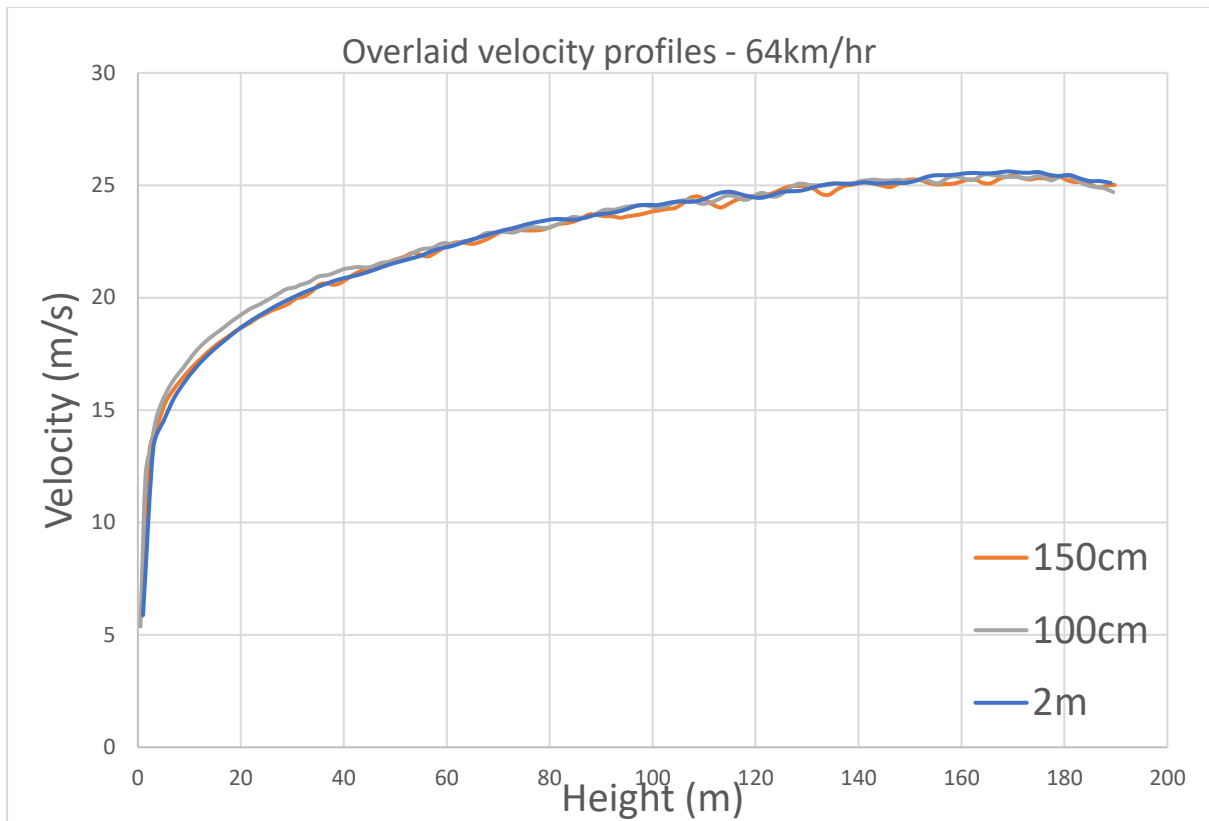


Figure 4-12. Velocity measurements close to the top surface of the domain were not included due to errors caused by flow parallel to an open boundary condition (as described in 4.3.3). These averaged wind velocity profiles were then used to calculate a CGI figure between the 2m and 1.5m cases, and the 1.5m and 1m cases.




Figure 4-12: Overlaid velocity profiles for 64km/hr cases, 150m from the inlet.

The relative error term ε was calculated as the average of the absolute value of the difference between each point in each profile, divided by the wind speed at the inlet. Velocities were interpolated when measurement positions did not match exactly, such as in 1.5m cases.

From the cases 150m from the inlet, the GCI values obtained were:

Table 4: GCI figures for various wind speeds.

37km/hr	$GCI_{2:1.5} = 14.080130$ $GCI_{1.5:1} = 4.012630$
46km/hr	$GCI_{2:1.5} = 8.485266$ $GCI_{1.5:1} = 5.65544$
54km/hr	$GCI_{2:1.5} = 8.594157$ $GCI_{1.5:1} = 4.419487$
64km/hr	$GCI_{2:1.5} = 27.221071$ $GCI_{1.5:1} = 3.4373801$

These GCI values between grid sizes of 1.5m and 1 m produce a GCI under 5 for wind flow in an open area in almost all cases. However, results from mid-canopy measurements show that

the presence of obstructions within the forested area led to a larger GCI, with particularly large errors in the sub-canopy region ($<40\text{m}$), as shown in Figure 4-13:

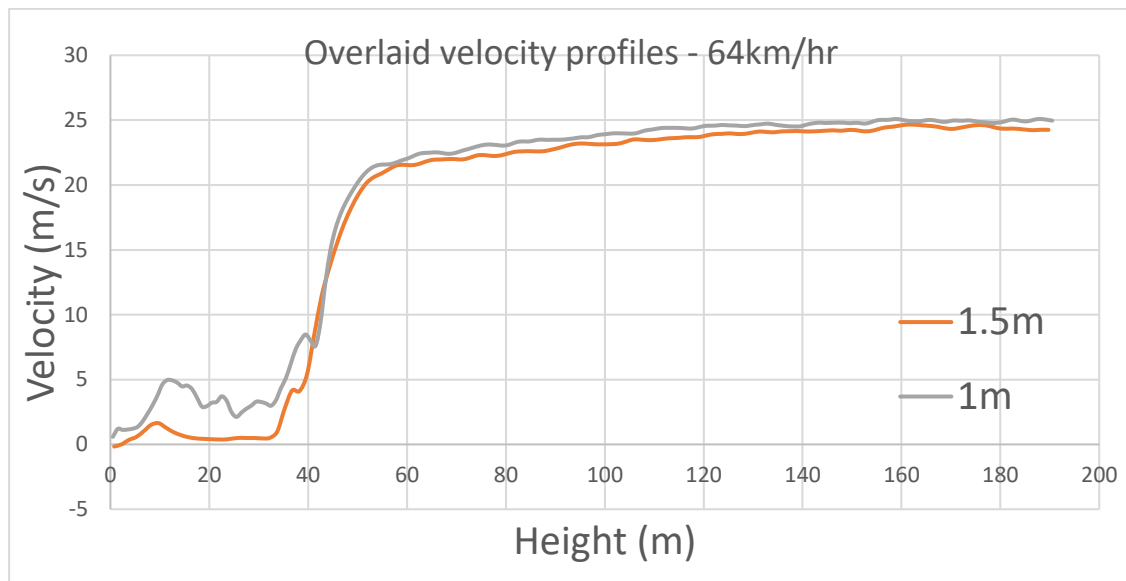


Figure 4-13: Overlaid velocity profiles for 64km/hr cases, at the centre of the forest.

These large discrepancies between the 1m and 1.5m cases in the sub-canopy area are due to the way that FDS models obstructions that only partially occlude a cell. In the 1.5m case, tree trunks that overlap multiple cells effectively obstruct a significantly larger volume than identically specified trunks in the 1m case. This causes erroneously low average sub-canopy wind speeds, despite flow in open areas being well-converged in the 1.5m cases. This effect is not present in 2m cases, as the specified tree dimensions fit fully within 2m grid cells. Due to this effect, 1m was chosen as the most appropriate grid size for our scenarios, as it is sufficient to produce a grid-converged wind field in open areas, while also avoiding errors in sub-canopy flow caused by multiple cells overlapping individual obstructions.

With 1m having been selected as the most appropriate grid size, some consideration must be given to the computational cost of running a scenario at this size. Attempting to simulate the whole domain at this grid size is not practical – we quickly run into limitations in terms of how much computational resources we can assign to each case. However, FDS allows us to subdivide the domain into a series of meshes, each of which can be given a separate grid size so long as that size is a multiple of two of any bordering mesh. As such, less important areas of the domain – where we don't expect any fires or the bulk of particle transport to occur – can be simulated at larger grid sizes to save computational resources. In addition to this, areas where we expect large amounts of activity to occur (such as in the area to be burned, and

immediately in front of it) can be split into a large number of computational meshes at 1m grid size to spread the computational load across a larger number of CPU cores.




Figure 4-14: Diagram of mesh layout throughout the domain, showing coarse and fine meshes near the inlet.

Using this feature of FDS to subdivide our domain to distribute computational load, the domain was split into 96 meshes (equivalent to two CPU sockets on the NCI Gadi supercomputer). 16 of these meshes are assigned to larger grid size regions (2m) along the sides of the domain, where we do not expect any fires and only expect some peripheral firebrand distribution to occur. The remaining 80 meshes are assigned to the central regions of the domain, split into eight equal-length sections in the x axis, five equal-width sections in the y axis, and two equal sections in the z axis. This subdivision of the domain produces relatively even computational time costs across each mesh, with the most computationally expensive mesh (mesh #33, the centre of the fire) only requiring approximately $\sim 3.6\times$ the compute time of the least expensive mesh (mesh #87, 2m grid size near the inlet).

4.5 Discussion of Assumptions and Limitations

Having now laid out the reasoning of how our simulation scenarios were specified (e.g. domain sizing, boundary conditions, atmospheric conditions, vegetation parameters) we should address the assumptions built into this set of scenarios and the limitations of the data that can be obtained as a consequence of these assumptions. Due to the limitations of computational resources available to us, we cannot simulate an arbitrarily large number of scenarios, so there are necessarily some compromises in which parameters we choose to vary and how much they are varied. These compromises place some limitations on the range of applicability of our model. For example, given that we have specified scenarios with fuel and atmospheric conditions that are extreme to catastrophic (FFDI 80 and above) we should not expect the resulting model to perform particularly well in low or medium FFDI conditions. Similarly, we have chosen a certain range of vegetation characteristics, heat release rates, and wind speed conditions for our simulations – wildfires that occur in forests with dramatically different conditions in these parameters may produce significantly different quantities and distributions of firebrands. However, while the conditions simulated in our set of scenarios may not be exactly like any given fire, the general statistical characteristics of the distributions and trends in distribution parameters obtained from our data are likely informative for a broad range of wildfire scenarios.

Some parameters that are likely to impact the distribution of firebrands have been kept constant across the range of scenarios – most notably, the height of the forest. While our selection of tree height was made with reference to relevant literature, there is clearly no single tree height that can be chosen as representative of all forests. The impact of this parameter on firebrand

distribution should be considered for future research, and the models produced from our data should only be considered as representative for forests with similar heights to those simulated in this data set (~40m).

4.6 Summary

By taking a methodical approach to the specification of physical parameters in our scenarios, we have produced a set of atmospheric, vegetation, flame and firebrand parameters that should allow us to properly simulate firebrand transport in conditions similar to that of real wildfire scenarios. The process by which we have chosen these parameters is thoroughly explained and based on published empirical studies of fire behaviour, firebrand generation, and studies of physical characteristics of firebrands.

The domain chosen for our scenarios has sufficiently large physical dimensions to capture all relevant fire, plume, and firebrand transport phenomena that we aim to investigate. These dimensions are 1200x500x200m. The effect of different boundary conditions on firebrand transport in the domain has been investigated and the most appropriate boundary conditions selected – open boundaries were selected for use at the inlet, top, and end of the domain, and mirror boundaries at the sides. A sensitivity analysis of domain height on wind velocity was conducted, showing no significant impact of an increase in domain height on wind velocity throughout most of the domain. A small but measurable reduction in wind velocity occurs close to the top surface of the domain due to the behaviour of flows moving parallel to an open boundary.

A grid convergence analysis was conducted using wind velocity profiles at a known distance from the inlet to determine convergence. This grid convergence analysis showed that a grid size of 1m was the most appropriate for our purposes. Due to the infeasibly high computational cost of simulating the whole domain at 1m grid size, the domain is subdivided into several regions. The central region of the domain, spanning 200m in width and running throughout the whole length and height of the domain, is simulated at 1m grid size. This 1m grid size region contains the area to be burned, and the vast majority of firebrand transport occurs in this region. The outer 150m on either side of this central 200m region is then simulated at 2m grid size, giving room for firebrands to move a large distance laterally.

Chapter 5 Simulation Results and Physics

Contents

Introduction.....	91
Simulation Output.....	91
Data Output	91
Plume Characteristics.....	92
Vortex Behaviour.....	99
Wake Vortices.....	101
Counter-Rotating Vortex Pairs	106
Particle Trajectory.....	112
Stream Tracers	112
Downwind Dispersion and Initial Height	119
Crosswind Dispersion	126
Overall particle distribution and trajectory	132

5.1 Introduction

Before investigating the statistical characteristics of firebrand distribution in our simulations, it is useful to investigate the distribution and transport of firebrand particles from a physics perspective. This chapter presents an overview of the data obtained from the set of simulations described in Chapter 4 before performing a deeper analysis of the physics of firebrand particle transport. Several transport phenomena are investigated, with particular focus on how the shape of the convective plume and vortices in the downwind region are impacted by the parameters of the simulation, and how these phenomena then influence the trajectory of firebrand particles. The trajectory of firebrand particles from their point of origin to final landing position can then be analysed both quantitatively and qualitatively with reference to the observed physics.

With the conditions of each scenario defined in the previous chapter, a final set of 46 simulation scenarios were simulated. This final set of simulations is the basis of most analyses going forward.

5.2 Simulation Output

5.2.1 Data Output

The output of an FDS simulation may include data on a wide range of physical characteristics of different components of the simulation, such as the motion, temperature, and chemical composition of solids and gases throughout the domain. To investigate the physical phenomena underlying particle transport in our set of simulations, it is necessary to process and interpret this data.

While a wide range of data on a variety of physical properties can be obtained From FDS simulations, not all data is necessarily relevant to the specific area we intend to investigate. For our purposes, we are primarily interested in data relating to the velocity and temperature of wind and data relating to the position and motion of particles. This data can be extracted in a range of different formats. FDS allows for data to be recorded as:

- Time series data at a particular point (analogous to a real measuring device, e.g. a thermometer)
- Time series of global data (such as overall HRR across the domain),
- Time series across a two-dimensional plane for fluid phase parameters (such as temperature and velocity)

- Three-dimensional data for fluid phase parameters at a particular moment in time across the entire domain
- Time series of data relating to particle position and physical characteristics.

For the set of simulations considered here, the primary data sets of interest are those that relate to the motion of the wind and particles moving through the domain. These data sets include:

- time series data of fluid phase characteristics for two-dimensional slices across the xy , xz , and yz planes,
- three-dimensional snapshots of wind velocity data at a particular instant in time, and
- time series data of particle positions.

The two-dimensional fluid phase time series data were recorded for the planes $y = 0$, $x = 400$ to $x = 600$ in 50m increments, and $z = 20$ to $z = 50$ in 10m increments. This data allows for detailed investigation of wind velocity in several locations throughout the domain, such as the path of the convective plume downwind of the fire, and the formation and motion of vortices beneath the plume.

Three-dimensional data was recorded for every cell in the domain every 10 seconds, starting at 10 seconds from the start of the simulation. This data set includes wind velocity and temperature data, allowing for detailed visualization of the wind velocity field at a range of times.

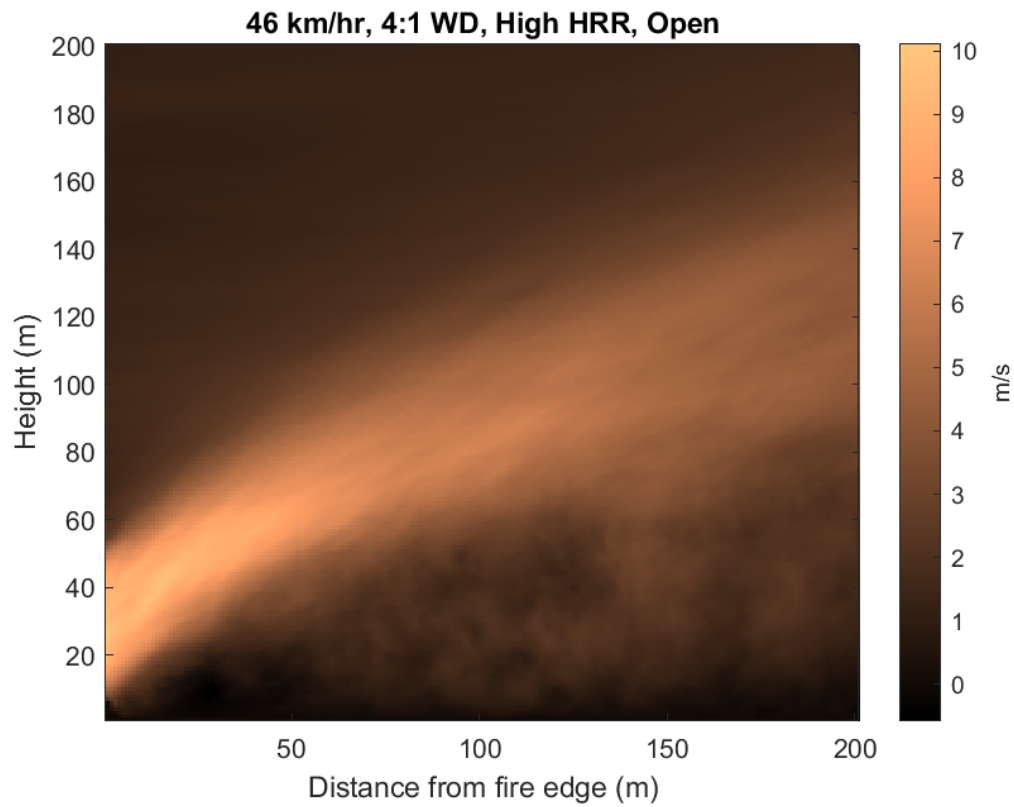
Both 2D and 3D data can be further analyzed to investigate more complex transport phenomena using in-house Python and Matlab code, as well as software tools like Paraview. For example, visualizations of time-averaged velocity components, streamlines, and measures of vorticity such as the Q criterion can be generated from these data sets.

Time series data relating to particle characteristics were recorded for each individual particle once per second from the moment each particle was introduced into the domain. This data records the position and temperature of the particle, allowing for the trajectory of each particle to be followed from insertion to final landing position.

5.3 Plume Characteristics

Simulation results can be visualized using Smokeview[73], a companion program for FDS that can display each simulation as a 2D or 3D animation. Each of the previously mentioned data outputs can be overlaid on this animation to produce a visualization of wind and particle

transport. Using these visualizations and data processing in MATLAB and Python, we can investigate and compare the shape of the plume against literature models of plume height. The most convenient way to visualize the shape of the convective plume is to produce a plot of the average vertical component of velocity through the plane $y = 0$, such as in Figure 5-1.



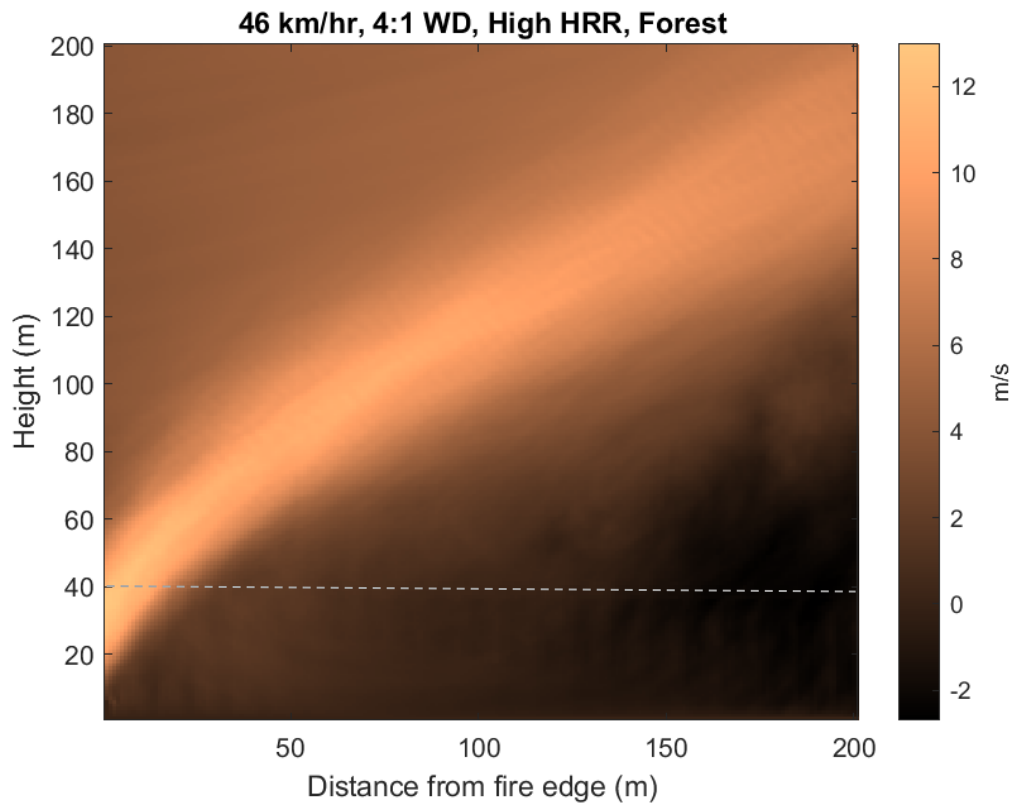


Figure 5-1: Plot of the average vertical component of velocity in the region immediately downwind of the fire. Velocity averaged over 55 seconds during the period of peak heat release rate. Dashed white line represents the top of the forest in the forested case.

We can observe from the plumes shown in Figure 5-1 that the plume initially exits the forest at a fairly steep angle. In the open field case, the plume is very well-defined within the first 50m of travel downwind, with clear edges between the plume and the surrounding area. Once the plume travels further downwind, it mixes with the air more thoroughly and the edges become less well-defined. In addition to this, there is a large volume of air underneath the plume getting entrained. The forested case retains a well-defined plume shape for a considerably longer distance, remaining relatively well-defined even 200m downwind of the edge of the fire. Very little air beneath the plume is getting entrained with it. The upward velocity of air in the plume is also faster in the forested case.

Using this visualization tool, we can check whether the convective plume is behaving as expected by comparing its shape to literature models. The Briggs plume-rise equation [74] is an analytical formula that describes the position of the center of a convective plume in a crosswind based on the heat flux and wind velocity. While this equation makes some assumptions about the shape of the heat source and plume that may not apply in our case, it is a good starting point for further analysis. The Briggs equation assumes that:

1. The Boussinesq approximation[75] applies
2. The plume has a vertical circular cross-section
3. There is a linear relationship between entrainment velocity and plume vertical velocity
4. Fluid properties are constant across this circular cross-section
5. Background wind velocity (U) is constant
6. Plume horizontal wind velocity is very close to U , so that the difference in velocity between the plume and background can be ignored
7. Background static stability is neutral ($N^2 = 0$)

The term N^2 is the square of the Brunt-Väisälä frequency [76], a term that describes the frequency that a hypothetical parcel of air would oscillate if it were displaced vertically by a small distance. This frequency can be calculated as:

$$N = \sqrt{\frac{|g|}{T_v} \frac{\Delta\theta_v}{\Delta z}} \quad (5-1)$$

Where g is gravitational acceleration, T_v is background virtual temperature, θ_v is virtual potential temperature and z is vertical position.

Experimental research has shown that plumes from wildfires tend to follow a shape similar to that predicted by the Briggs equation[77]. In addition to this, many other models of plume shape are derived from the Briggs equation. The Briggs plume-rise equation is:

$$z = \left[\left(\frac{3}{2\beta^2} \right) \frac{B_{flux}}{\pi U^3} \right]^{\frac{1}{3}} x^{\frac{2}{3}} \quad (5-2)$$

Where z is plume height, x is the distance in the downwind direction, U is wind velocity, B_{flux} is the buoyancy flux, and β is a constant (0.6).

Several modifications of the Briggs equation are described by Fisher et. al. [78] and were tested against data obtained from a set of experimental pool fires. These modified equations can be expressed as an equation of the form:

$$\log_{10}z(x) = a\log_{10}x + b\log_{10}Q_E - c\log_{10}U + d \quad (5-3)$$

Where the terms a , b , and c are dimensionless coefficients, and d has units such that the dimensions of the result are correct. Q_E is the non-radiant heat release rate, approximated as ~70% of the total HRR in literature. The assumption that 30% of heat is radiated is a reasonably good approximation for our simulations as the FDS HRR output files show that during the burning period, between 20-27% of heat released is emitted as radiant heat. Of the plume rise

models tested by Fisher et al [78], the Zonato et. al. model [79] was the best performing model. Using the Zonato et. al. model of plume height from free-burning fires, we can compare the shape of the plume obtained from our simulations with an experimentally tested plume model. This is shown in Figure 5-2:

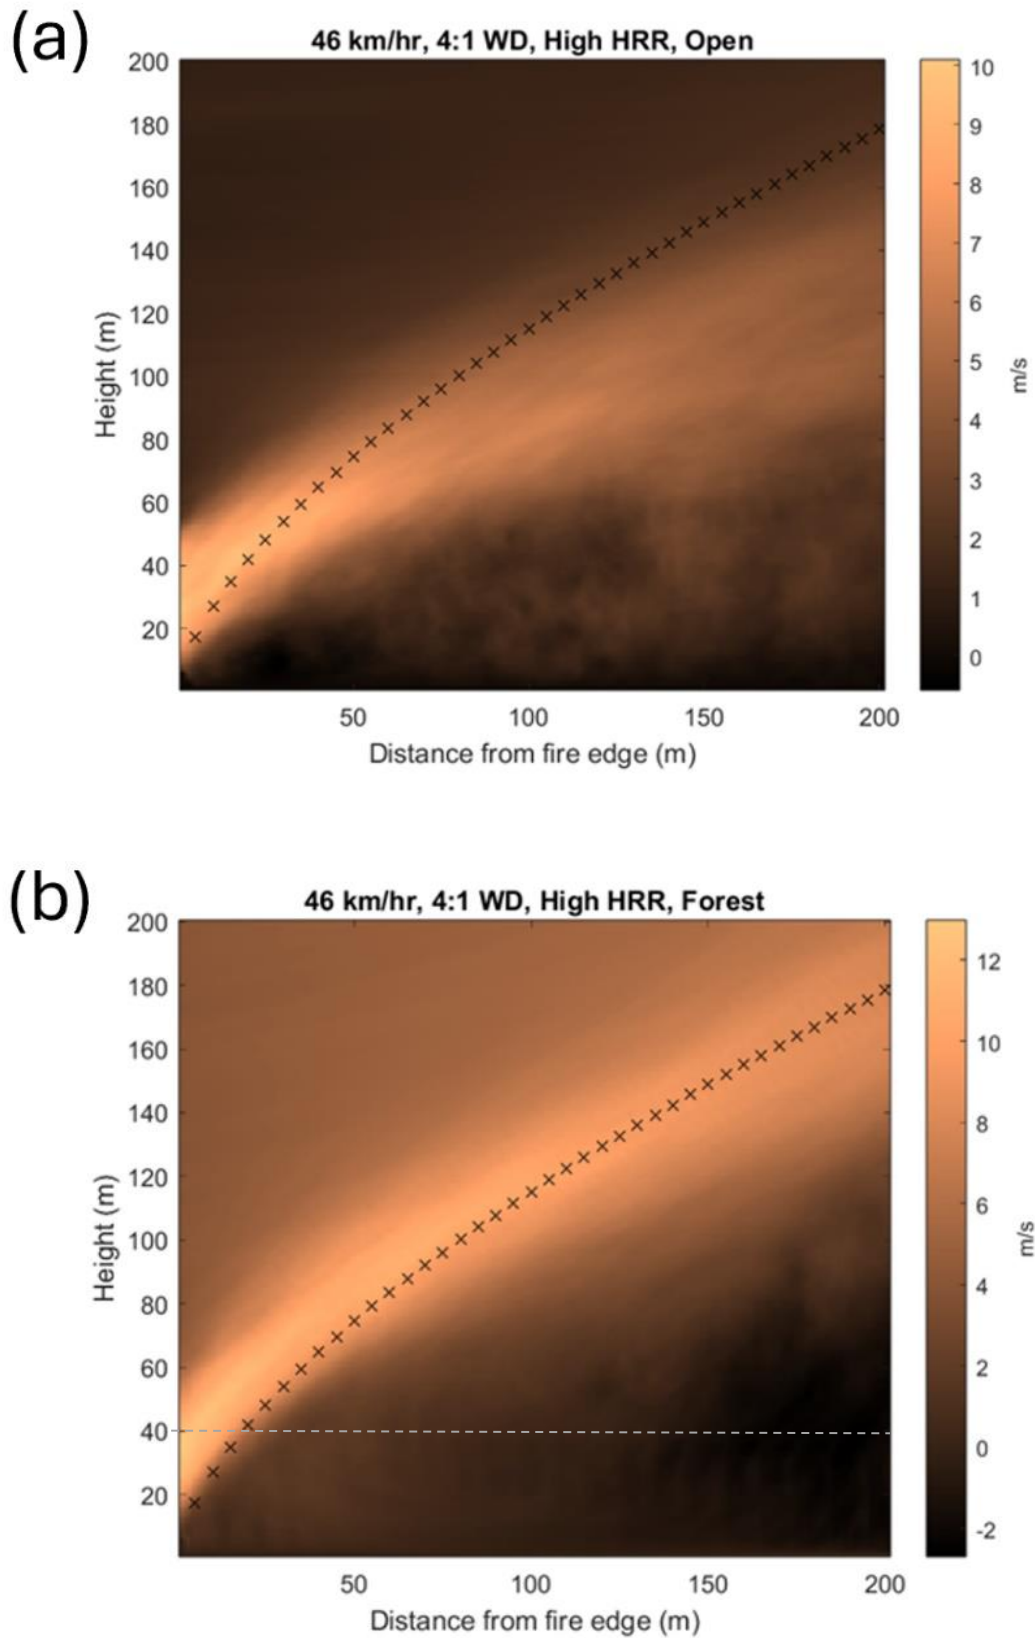


Figure 5-2: Simulated plumes with literature Zonato model of plume rise overlaid. Dashed white line indicates forest height in the forested case.

The Zonato model of plume rise predicts the path of the simulated plume quite well in most circumstances. The model performs better for simulations that have forested domains and for simulations with physically larger fires and higher heat release rates. In cases with open fields downwind of the fire, the plume usually does not rise as high as the Zonato model would predict. This may be because some of the energy that would otherwise contribute to plume rise is instead used in the creation of vortices close to the ground, producing the large volume of rising air underneath the plume visible in Figure 5-2. The presence of a forest below the plume inhibits the formation of vortices near the ground. The continuity of the forest may also play a role in the difference in plume behaviour between forested and open field cases – the sharp edge of the forest in the open field cases may cause a disturbance in the wind velocity field beneath the plume and result in a different plume shape.

The Zonato model performs poorly for the very small and low intensity cases. In these cases, the plume rises more steeply than the Zonato model would predict. This can be seen in Figure 5-3. It seems that this effect is caused by the forest shielding the plume from ambient wind for the first few seconds of travel – this keeps the plume traveling almost vertically until it gets well above the canopy. Once the plume reaches greater heights, it begins to flatten and follow a path almost parallel to the Zonato model prediction. These observations about the applicability of the Zonato model in the context of wildfires may be novel – particularly in how they relate to changes in plume behaviour near the edges of forests and in certain lower-intensity fire scenarios.

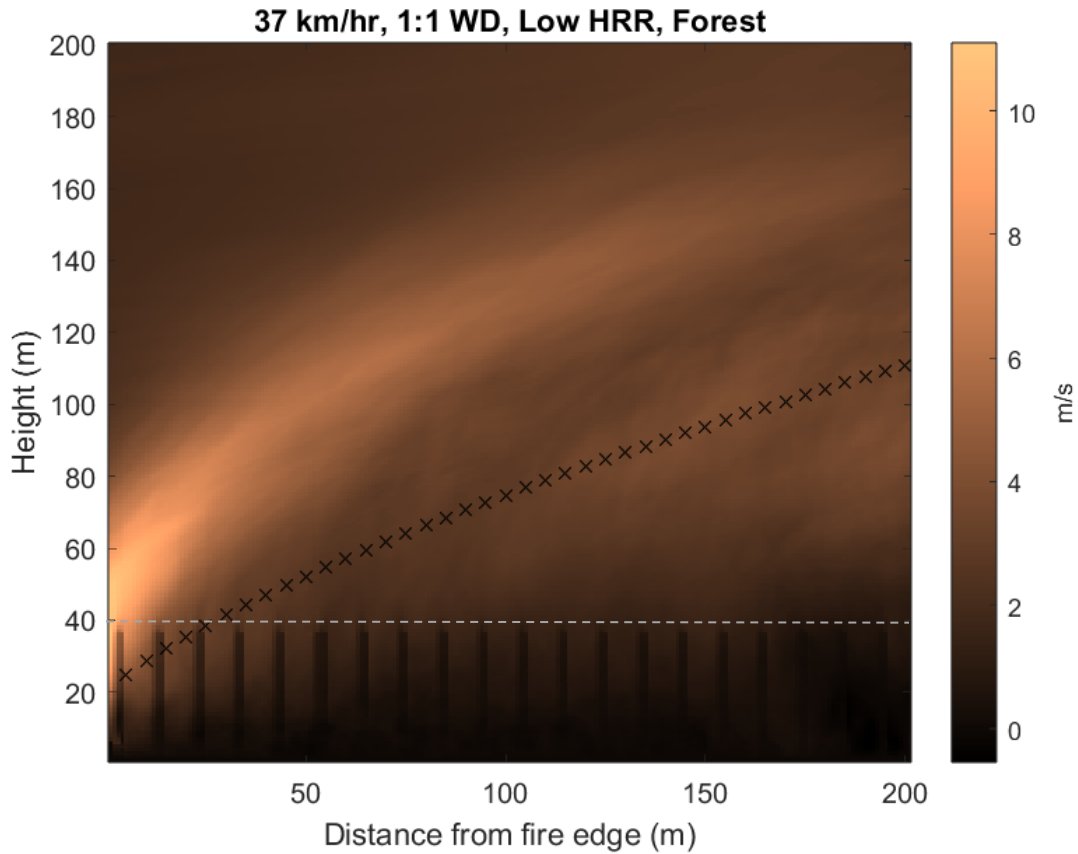


Figure 5-3: Zonato model of plume rise overlaid on a simulated plume from a small, low-HRR fire case. The plume rises higher than the Zonato model would predict.

5.4 Vortex Behaviour

The formation of vortices in the region downwind of a fire or heat source has been observed in numerous real wildfire scenarios [80, 81], field-scale experimental fires [82], laboratory-scale experiments [83, 84] and simulations [62, 85]. Vortices have been observed to form at a wide range of length scales and may be oriented in different directions depending on the position, shape, and size of the fire and the atmospheric conditions present. Vortices may either move far from their origin or remain relatively fixed in position depending on the type of vortex. As described in Lareau et. al., “embedded vortices” are those that remain anchored to the fire, and “shedding vortices” are those that detach and travel downstream[86]. Figure 5-4 shows a number of common vortex types downwind of a wildfire, giving an example of both embedded (e.g. CVP) and shedding vortices (e.g. wake vortices).

Given the body of literature regarding vortex formation around wildfires and other heat sources, we should expect to observe the formation of vortices in the region downwind of the fire and beneath the convective plume across the range of scenarios being studied here. These

vortices will likely have a significant effect on the trajectory and distribution of firebrand particles as they travel downwind. By visualizing the location of these vortices and quantifying the wind velocity throughout their volume, we can determine how the different initial conditions of each scenario may impact particle transport from a physics perspective.

In literature, several types of vortex downwind of a fire are commonly identified [80, 87]. The most prominent vortex formation is usually a large Counter-rotating Vortex Pair (CVP) that forms underneath the plume. These counter-rotating vortices rotate around an axis parallel to the direction of plume rise. In our case, this means that the counter-rotating vortex pair is primarily rotating around the X-axis, and would be most visible in slices of the domain that show a Y-Z Plane. Beneath the counter-rotating vortex pair, wake vortices are formed. These vortices rotate around the vertical axis while traveling downwind, and as such are most easily viewed in the X-Y plane. Finally, small “wall vortices” that travel close to the ground are formed. Due to the small size of these vortices and their location very near the ground, these vortices are unlikely to impact the overall distribution of firebrand particles significantly. Additionally, the small size of these vortices makes them difficult to capture at the grid resolutions used in these simulations – refining the grid size close to the ground sufficiently to capture these very small vortices would be prohibitively computationally expensive. Moreover, it is out of the scope of this study




Figure 5-4: Diagram of vortices formed beneath a convective plume in a crosswind. Source: Shinohara and Matsushima, [83]

5.4.1 Wake Vortices

Wake vortices are a type of vertically-oriented vortex that tend to form close to the downwind edge of a fire. They are produced in pairs with opposite rotation, and are a type of shedding vortex – they tend to travel downwind rather than remain attached to the fire. These vortices have been observed in wildfire conditions and large-scale experiments, such as the 1000MW Météotron experimental fires by Church et al [88] - an example of a pair of wake vortices photographed during these experiments is shown in Figure 5-5.

These vortices are likely to have a significant impact on the distribution of firebrand particles due to their relatively large size and horizontal velocity components. Wake vortices typically originate near ground level and rise several meters vertically upwards, meaning these vortices may have a greater impact on the trajectory of particles that travel in shallow arcs than particles that are lofted to high altitudes.



Figure 5-5: Paired wake vortices on either side of a fuel burner heat source. The vortices are rotating in opposite directions – the vortex on the left showing cyclonic rotation and the vortex on the right showing anticyclonic rotation. Source: Church, Snow and Dessens [82]

As with the shape of the plume, we can observe the behaviour of wake vortices by analyzing wind velocity data. Using visualization tools such as Smokeview and by post-processing velocity data with Python, we can observe the formation of vortices in the wind velocity field and see how those vortices may impact the distribution of firebrand particles. For this section concerning wake vortices, the wind velocity field is visualized using Smokeview and particle distribution is visualized in Python. Smokeview is used to render the time series data of wind velocity over a plane as an interactive animation, so that the movement of vortices across the domain can be observed visually. Python is used to produce kernel density estimate plots of particle distribution from raw data, such as Figure 5-8.

Unlike the previous section where convective plumes were visualized using a time-averaging approach, we cannot use a time-averaging approach to visualize wake vortices easily. While the plume and counter-rotating vortex pairs tend to remain fixed in position, wake vortices tend to migrate downwind over time. These moving wake vortices can be visualized at a particular instant fairly easily, but are difficult to track and characterize over a long period. The movement of these vortices can be seen by comparing the 2D visualization of the lateral velocity (v -velocity) vectors through a planar slice at 50m altitude at two different points in time and finding paired regions of positive and negative v -velocity, as in Figure 5-6.

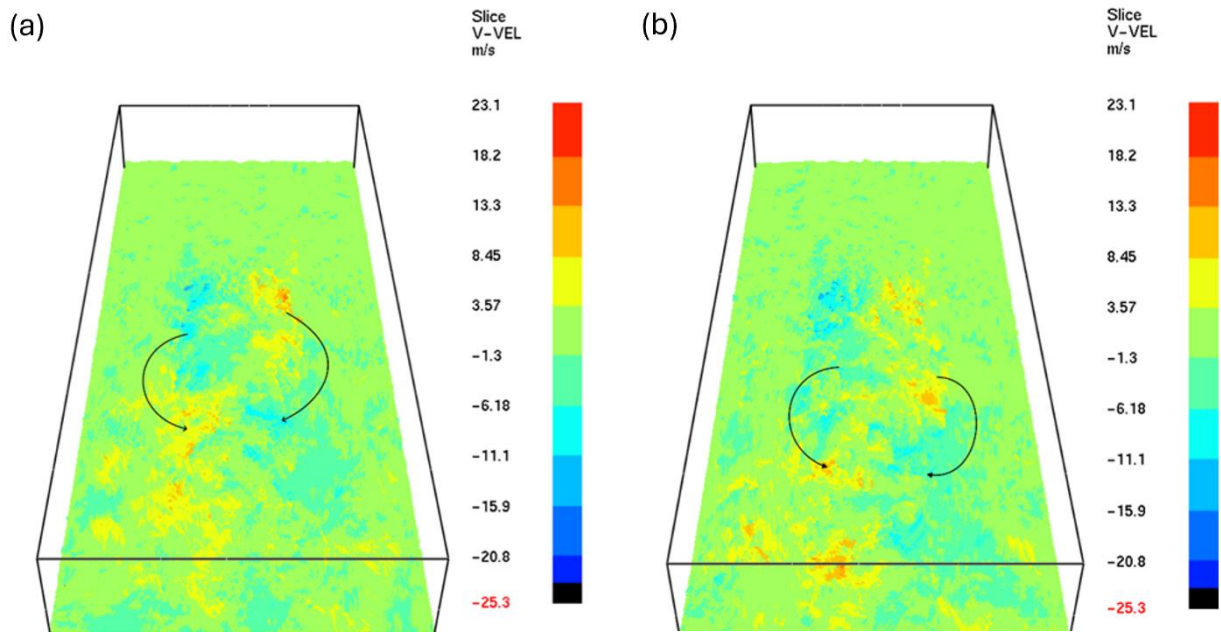


Figure 5-6: Example of wake vortices migrating downwind in the 54 km/hr, 2:1WD, High HRR, Open Field case. Vortex direction is indicated by arrows. (a): Visualization at 92.6s, (b): Visualization at 102.6s.

In forested domains, these wake vortices are almost completely eliminated in the region beneath the canopy. This can be observed by comparing v -velocity data from the slice at 30m altitude between a comparable forested and non-forested case, such as in Figure 5-7a and Figure 5-7b. In this diagram, the colour bar has been modified to show a histogram of cell v -velocity, clearly indicating the effect of the forest on the formation of vortices in the sub-canopy region – the forested cases (Figure 5-7b, d, f) show a very narrow spread of v -velocity, with most of the plane showing values close to zero. The open field cases (Figure 5-7a, c, e) show a wider spread of velocities with long tails, highlighting cells around the leading and trailing edges of the wake vortices where v -velocity is highest.

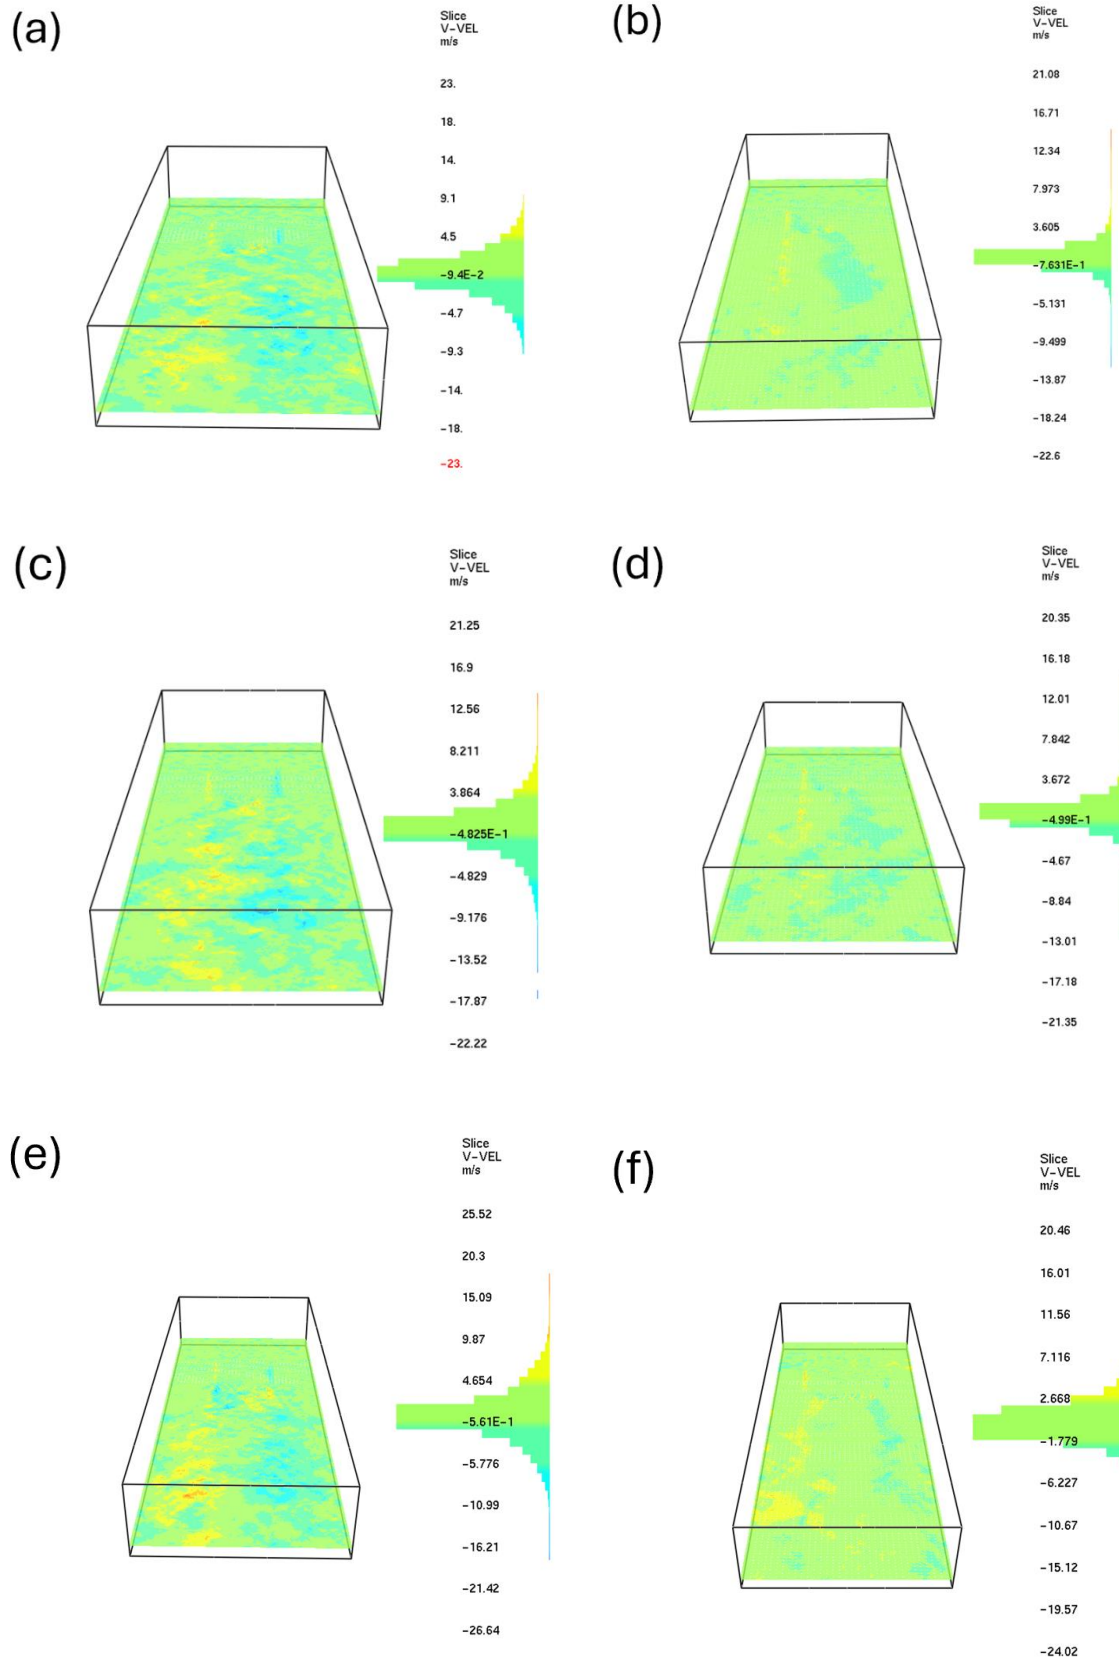


Figure 5-7: Planes of v-velocity for open field and forested cases ($Z=30\text{m}$). (a) Data from the 54 km/hr, 2:1WD, High HRR, Open Field case. (b) 54 km/hr, 2:1WD, High HRR, Forested case. (c) 54 km/hr, 2:1WD, Low HRR, Open Field case. (d) 54 km/hr, 2:1WD, Low HRR, Forested

case. (e) 54 km/hr, 4:1WD, High HRR, Open Field case (f) 54 km/hr, 4:1WD, High HRR, Forested case

Comparing cases with different HRR values (e.g. Figure 5-7a and 7c), we can see that reducing HRR has a noticeable impact on the histogram of v -velocities at 30m altitude. The distribution of v -velocity becomes slightly narrower, with a greater proportion of cells in this plane showing velocities close to zero. This is indicative of weaker and smaller wake vortices formed by a correspondingly less intense fire. Similarly, as W:D ratio increases (and thus the area of the burning region increases) the histogram of v -velocities widens – all else being equal, a larger fire will produce larger and more intense wake vortices.

The lack of strong wake vortices in the sub-canopy region of forested scenarios may be a contributing factor to the observed difference in crosswind firebrand distribution between forested and open field cases. The crosswind distribution of particles in forested domains is observed to be much more uniform than in open field domains with similar conditions. This difference in crosswind distribution can be observed clearly in many cases, and is particularly noticeable in cases with larger W:D ratios. For example; the 54 km/hr, 4:1 WD, High HRR cases (Figure 5-8a, 8b) show a dramatic difference in the shape of the crosswind marginal distribution, with the forested case (Figure 5-8a) showing a shape that is relatively flat compared to the strongly two-peaked shape of the open field case (Figure 5-8b).

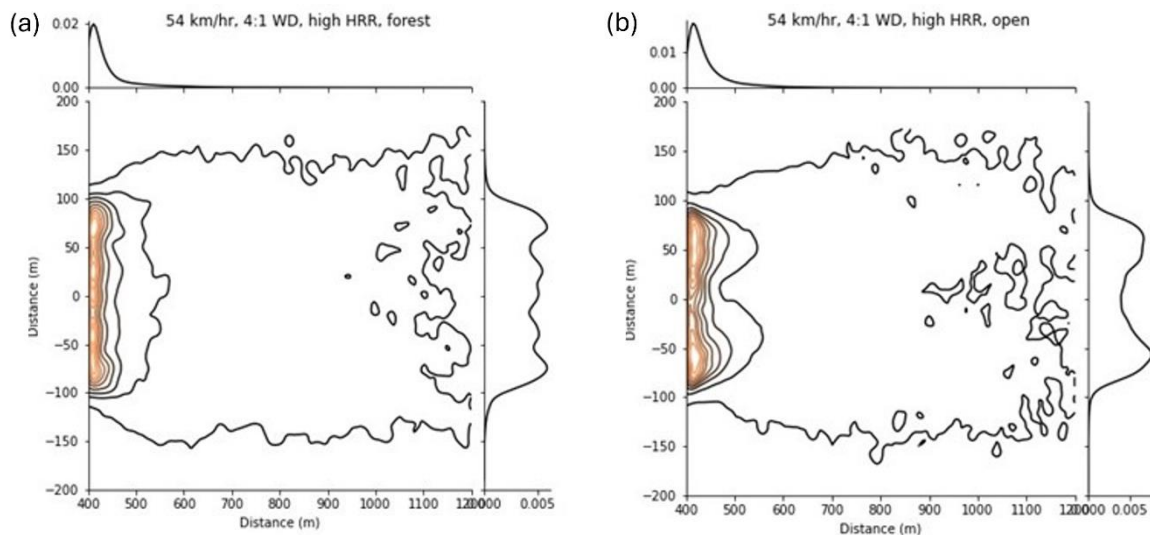



Figure 5-8: Comparison of particle distribution between comparable forested and open field scenarios using KDE. The forested case shows a much more uniform distribution of particles in the crosswind marginal distribution.

5.4.2 Counter-Rotating Vortex Pairs

The formation of a counter-rotating vortex pair in the region beneath a convective plume or jet in a crossflow is widely reported in literature, with counter-rotating vortex pair (CVP) formation observed in naturally occurring wildfires [86], laboratory-scale wind tunnel experiments [80, 89], and in simulations of airflow around a fire [87]. As described previously, these vortices are “embedded” – meaning they remain anchored to the heat source that they originate from. Given that average wind speed at the inlet of the domain is held constant in each scenario and the fire is specified with a fixed heat release rate, the counter-rotating vortex pair should have fairly consistent properties over time once it is established. The position of these vortices beneath the convective plume and close to the fire means that they will likely play a significant role in the transport and distribution of firebrand particles.

Due to the relatively static position of the counter-rotating vortex pair underneath the plume, it is fairly simple to produce a time-averaged visualization of these vortices and to produce quantitative data about the size, location, and velocity components of these vortices. These characteristics allow for a more detailed analysis of the impact of different scenario conditions on the behaviour of these vortices. Plotting the average v -velocity and w -velocity through a 200x200m section of the $x = 500\text{m}$ slice, we can obtain images of the counter-rotating vortex pair as shown in Figure 5-9.



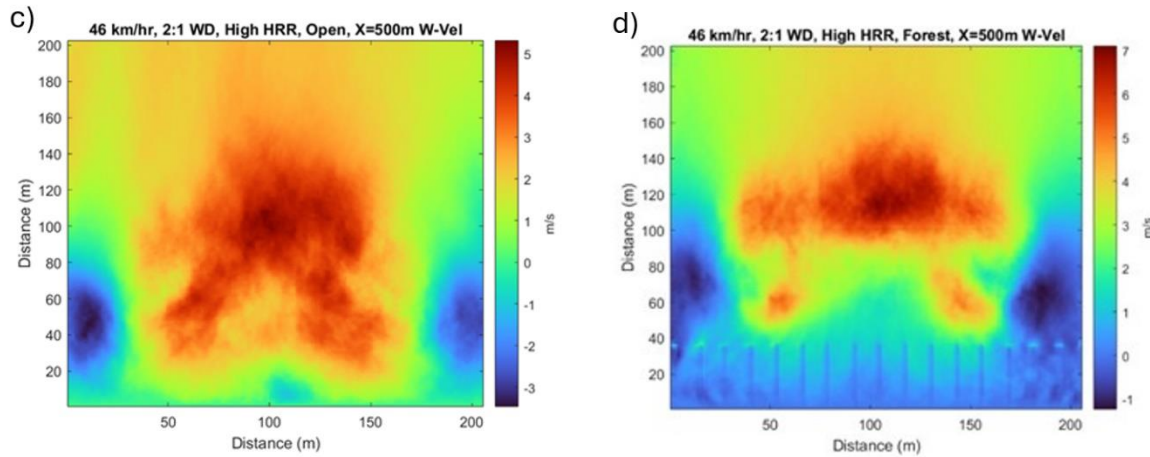


Figure 5-9: Visualization of average v -velocity and w -velocity through the slice $X=500\text{m}$ for the 46km/hr , $2:1\text{ WD}$, High HRR cases.

We can identify from these visualizations that the counter-rotating vortex pairs are significantly affected by the forest and canopy. In the forested case, the paired regions of positive and negative v -velocity that represent the upper and lower regions of each vortex are significantly closer to each other, producing a smaller and faster vortex than in the open field case. The regions of maximum and minimum v -velocity are shifted upward in the forested case – the fastest-moving areas of the vortex pair are moved away from the ground and toward the top of the canopy. The w -velocity plots show that the forested case produces a smaller but much faster region of upward-moving air between the vortices.

While the difference between the velocity fields in forested and non-forested cases can be observed visually, a quantitative analysis may allow for greater insight. We can measure the average absolute value of v -velocity and w -velocity components across the slice to obtain a quantitative measure of both vortex strength (v -velocity) and plume rise (w -velocity). We can also approximate the location of the vortex centres by identifying the location of maxima and minima in v -velocity in the left and right halves of the slice – the centre of the vortex should lie on a line between these points where v -velocity is zero. This simple quantitative analysis reveals the magnitude of the impact of different initial parameters on the strength of the vortices and plume rise in the $x = 500\text{m}$ slice.

The most significant factor in vortex behaviour is the presence of a forest in the downwind region of the domain. In scenarios where the domain was forested, the average v -velocity across the slice was reduced by 18%, but the average w -velocity was increased by 45% - a significant difference in behaviour. This change in v -velocity magnitude relative to w -velocity in forested cases is likely to explain some of the difference in particle distribution patterns between open

field and forested cases. In general, the wider forested cases tend to show a more uniform distribution of particles in the crosswind direction than open field cases – as shown in Figure 5-8. This may be because the slower v -velocity component of vortices in forested conditions is less able to scatter particles laterally, and particles lofted directly through the center of the domain into the faster upward w -velocity regions are carried further. The position of the centre of the counter-rotating vortex pair was also shifted upward by an average of 54% relative to the open field cases, but there was no consistent effect on the position of the vortices in the horizontal axis. This shift in the position and behaviour of the CVP in forested domains is particularly interesting, and should be noted. Although counter-rotating vortices have been studied extensively, there may be novel findings about their behaviour at low altitudes interacting with forest canopies (especially where the forest canopy may be interrupted – e.g. at a firebreak).

The overall effect of HRR on vortex behaviour seems to be a uniform but relatively small increase in both w -velocity and v -velocity. The high HRR cases, with 50% more HRR per unit area, produced vortices with only 12% greater v -velocity and 13% greater w -velocity on average. The increase in w -velocity reflects both faster vortex rotation speed and an increase in the strength of the convective plume. The effect of HRR on the position of the vortex centers is relatively small, with a negligible effect on their position in the horizontal axis and a small increase in vertical position for open field cases.

The effect of the W:D ratio on vortex speed is a general upward trend: larger fires tend to have greater average v -velocity. The rate of v -velocity increase has strongly diminishing returns as the W:D ratio increases. Cases with 2:1 W:D ratios produce roughly 30% faster v -velocity compared to 1:1, but 4:1 WD cases are not consistently faster than 2:1. The effect is similar for w -velocity: the 2:1 cases have approximately 40% faster w -velocity than 1:1, but the 4:1 cases are not consistently faster than 2:1. The position of the vortices is strongly impacted by the width of the fire, as expected – wider fires (width being affected by both wind speed and W:D) tend to produce vortices centered further apart in the horizontal y axis. However, the effect of flame width on vortex y -position is not linear, as shown in Figure 5-10.




Figure 5-10: Plot of flame edge y-position vs. vortex y-position

We can compare the ratio of the y-position of the edges of the fire (y_f) and the y-position of the centres of the vortices (y_v) to observe this nonlinear trend more clearly, as in Figure 5-11. Plotting the ratio of y_v/y_f across y_f , we can see that the trend could be roughly approximated as a function of the form $\frac{y_v}{y_f} = Ae^{y_f+B}$ ($r^2 = 0.643$). The narrowest fires produce vortices with y-coordinates that lie up to 2.6 times further from the line $y = 0$ than the flame edge, while the largest fires tend to produce vortices that are almost exactly in line with the flame edge ($y_v/y_f = 1$).




Figure 5-11: Plot of y_v/y_f across y_f at $x=500\text{m}$. A trendline showing a power law relationship is shown.

The tendency of the narrower fires to produce vortices that spread far outwards from the flame edge relative to the width of the flame may be responsible for some of the differences in the shape of the particle distribution contours between the wider and narrower fire scenarios. By comparing the distribution of particles between cases with a large y_v/y_f ratio and cases with smaller y_v/y_f ratios and marking the position of the vortices at $x = 500\text{m}$, we may be able to see what effect vortex position has on the distribution. Given that the presence of a forest in the domain is known to affect the vertical position and size of the vortices, we will compare both forested and non-forested cases.

Figure 5-12 shows kernel density estimate plots of particle landing position overlaid with the position of the vortices along the $x = 500\text{m}$ line for a selection of 37 km/hr cases with a wide range of y_v/y_f values. The forested 1:1 WD case (Figure 5-12a) has the largest y_v/y_f of all cases shown.

The position of the vortex centres tends to lie just outside the second least dense probability density contour. Comparing the forested 1:1 WD (Figure 5-12a) case to the open field 1:1 WD case (Figure 5-12c), we can observe a wider dispersion of particles throughout the domain. This may be due to the greater separation between the vortices, pulling particles further from

the centre line and dispersing them across a greater distance. However, the impact of the y_v/y_f ratio on the overall shape of the distribution is difficult to separate from other factors – it is hard to definitively say that the greater width of the distribution in the 1:1 WD forested case is due to the greater separation between the vortices, or that the wider distribution and greater vortex separation are both caused by a combination of other factors. The two cases with 4:1 WD (Figure 5-12b, d) have very similar y_v/y_f ratios but produce a noticeably different distribution of particles due to the different forest conditions in each case.

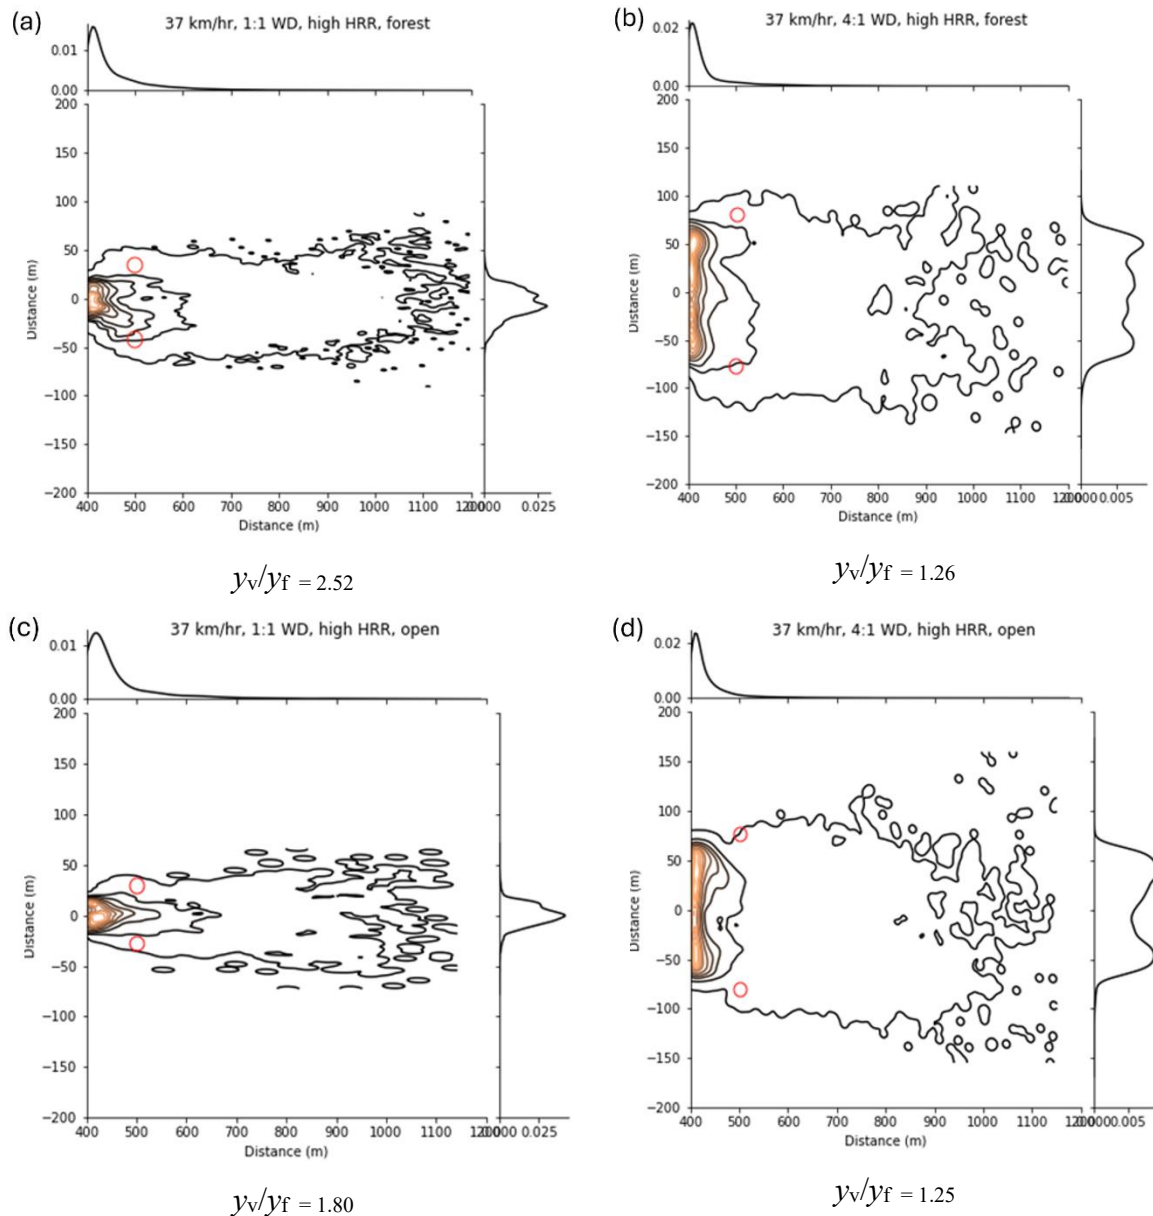


Figure 5-12: Kernel density estimate plots of particle landing density with vortex positions along the $x = 500\text{m}$ line circled.

5.5 Particle Trajectory

Having described how each scenario parameter impacts the motion and position of vortices in the wind velocity field, we may investigate how these parameters and vortex behaviours impact the characteristics of particle trajectories. To understand the effect of each parameter on particle trajectory, we will first take a qualitative approach and examine the shape of the streamlines throughout the domain. This will allow us to make some general observations about how a particle inserted under different initial conditions in a given area of the domain may move. Then, we may take a quantitative approach and try to describe the impact of simulation parameters on the shape of particle trajectories – looking at quantitative measures like the peak particle altitude each species of firebrand obtains under different conditions, the effect of particle insertion height on final landing distance, and the lateral spread of particles from their initial position in the crosswind axis.

5.5.1 Stream Tracers

To illustrate some general trends in the shape of the path of massless tracers, six cases have been selected for use as examples. These cases are:

1. 46 km/hr, 2:1 WD, High HRR, Forested (Figure 5-13)
2. 46 km/hr, 2:1 WD, High HRR, Open Field (Figure 5-14)
3. 64 km/hr, 2:1 WD, High HRR, Forested (Figure 5-15)
4. 64 km/hr, 2:1 WD, High HRR, Open Field (Figure 5-16)
5. 54 km/hr, 4:1 WD, Low HRR, Forested (Figure 5-17)
6. 54 km/hr, 4:1 WD, Low HRR, Open Field (Figure 5-18)

Each of these cases is shown from two perspectives using Paraview visualization software: a side-on view showing the X-Z plane, and a top-down view showing the X-Y plane. The stream tracers are inserted uniformly throughout a 120m radius sphere centered at the downwind edge of the fire ($x = 400\text{m}$) at an altitude of 60m. The velocity field used for this visualization is a snapshot of the domain at 100s, during the period of peak flame intensity and particle insertion. The path of the massless tracers is coloured by speed.

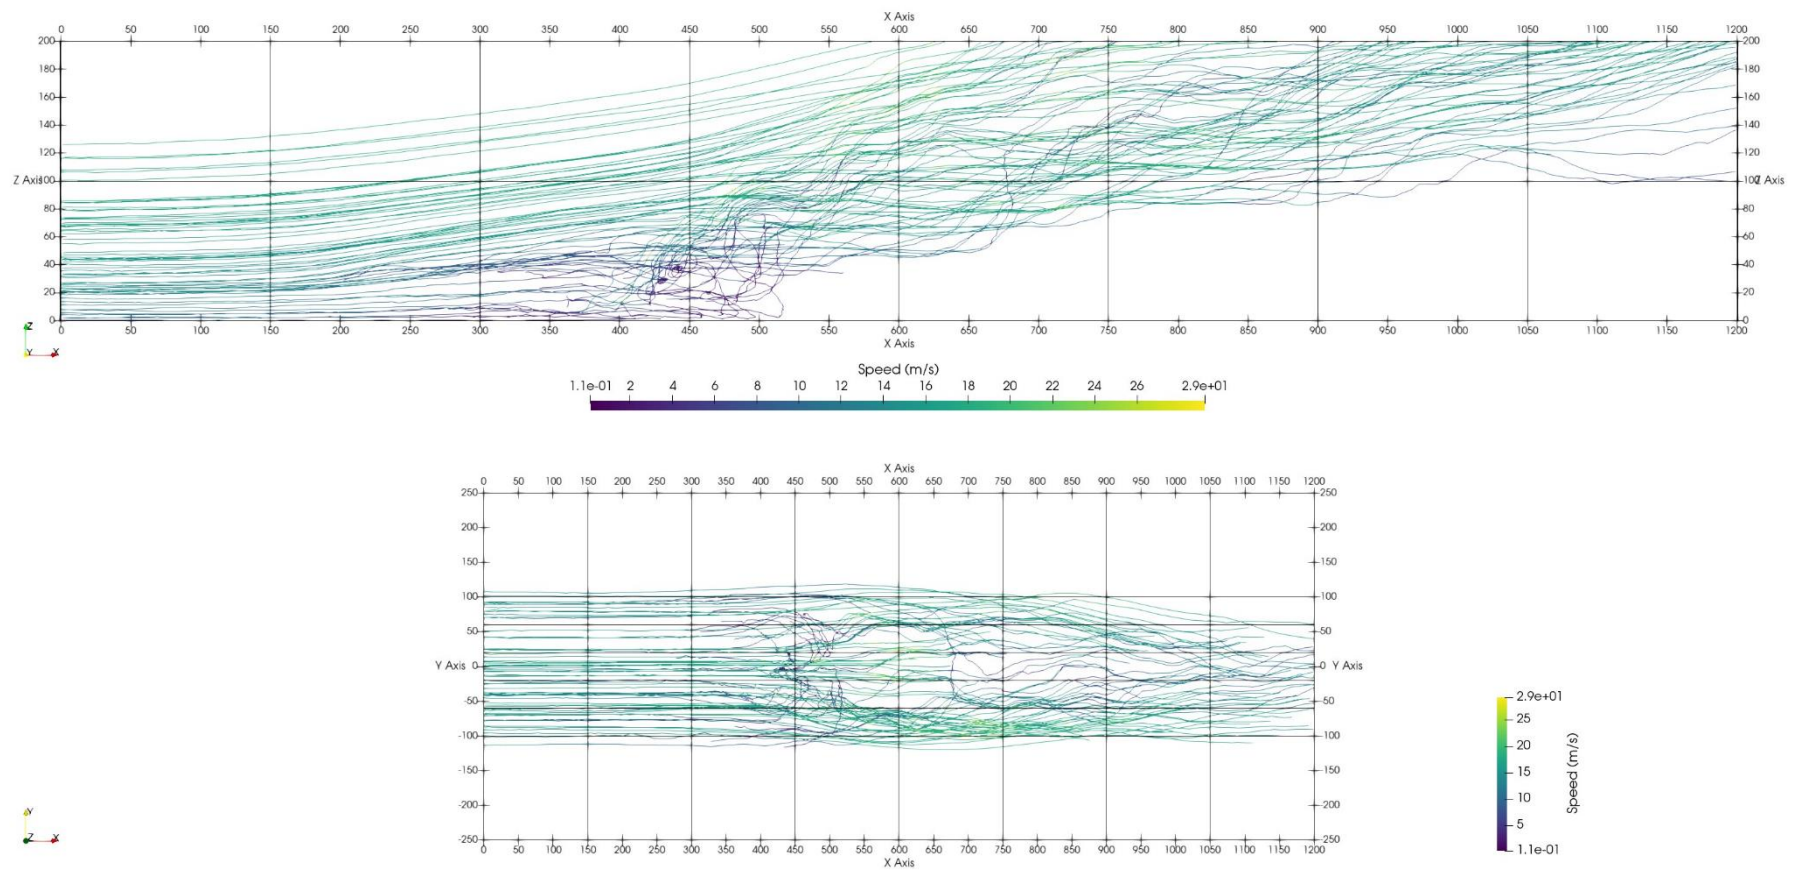


Figure 5-13: Stream tracers for the 46km/hr, 2:1 W:D, High HRR, Forested case. Top: View in the X-Z plane. Bottom: View in the X-Y plane.




Figure 5-14: Stream tracers for the 46km/hr, 2:1 W:D, High HRR, Open Field case. Top: View in the X-Z plane. Bottom: View in the X-Y plane.




Figure 5-15: Stream tracers for the 64km/hr, 2:1 W:D, High HRR, Forested case. Top: View in the X-Z plane. Bottom: View in the X-Y plane.




Figure 5-16: Stream tracers for the 64km/hr, 2:1 W:D, High HRR, Open Field case. Top: View in the X-Z plane. Bottom: View in the X-Y plane.

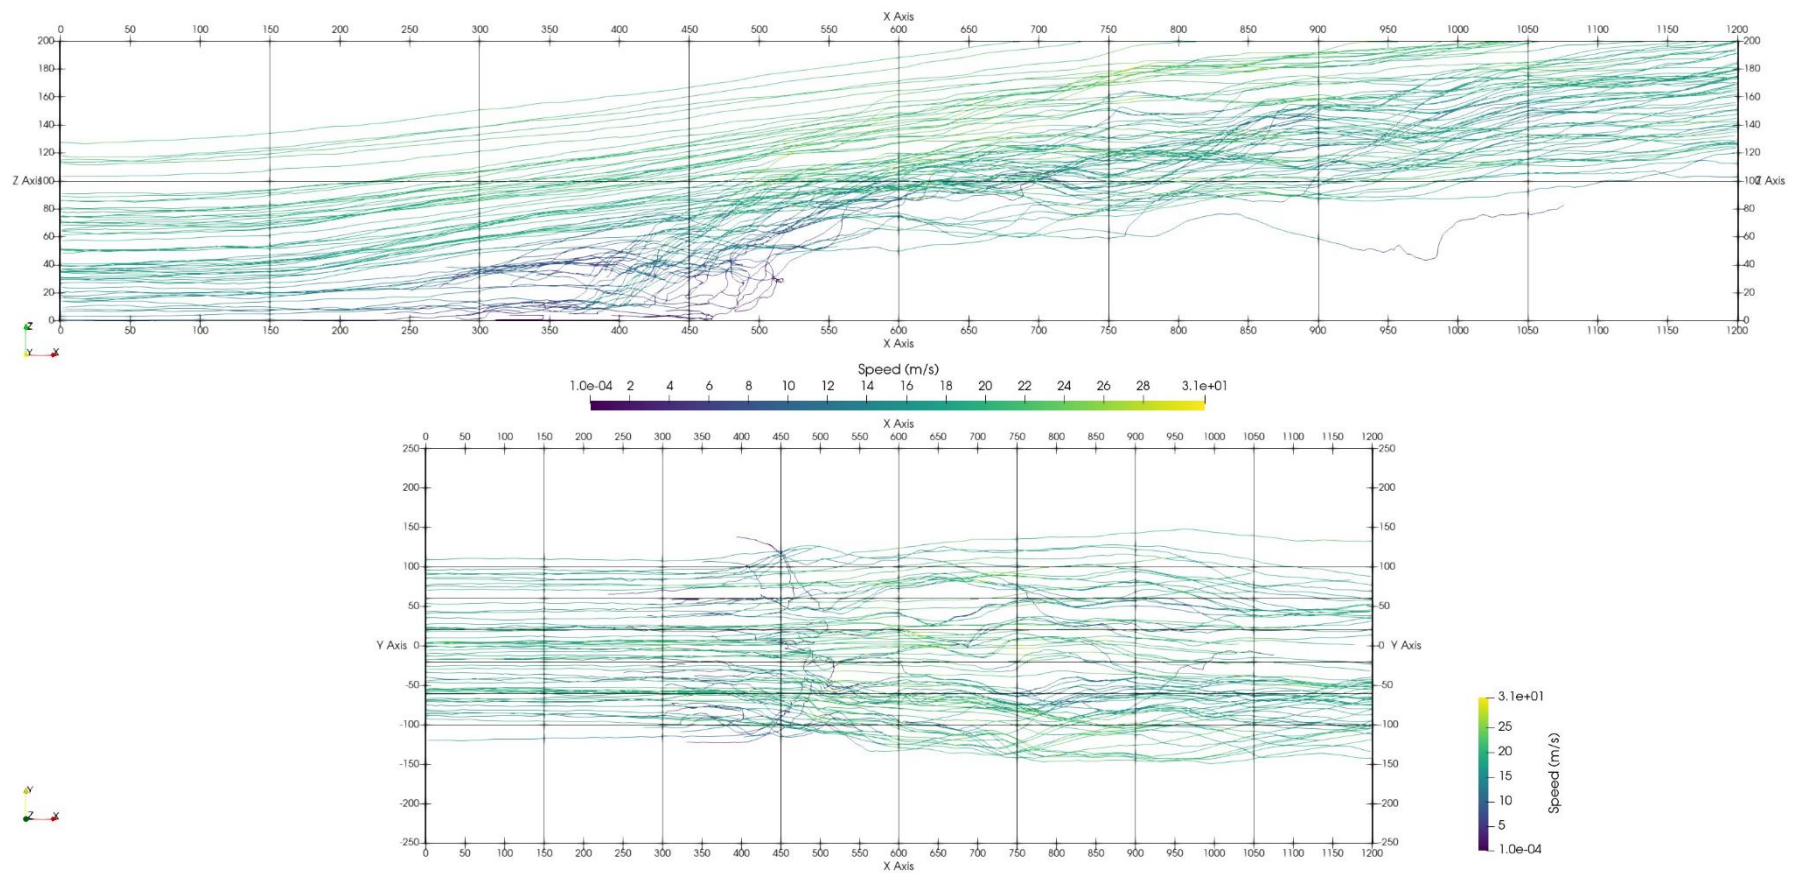


Figure 5-17: Stream tracers for the 54km/hr, 4:1 W:D, Low HRR, Forested case. Top: View in the X-Z plane. Bottom: View in the X-Y plane.

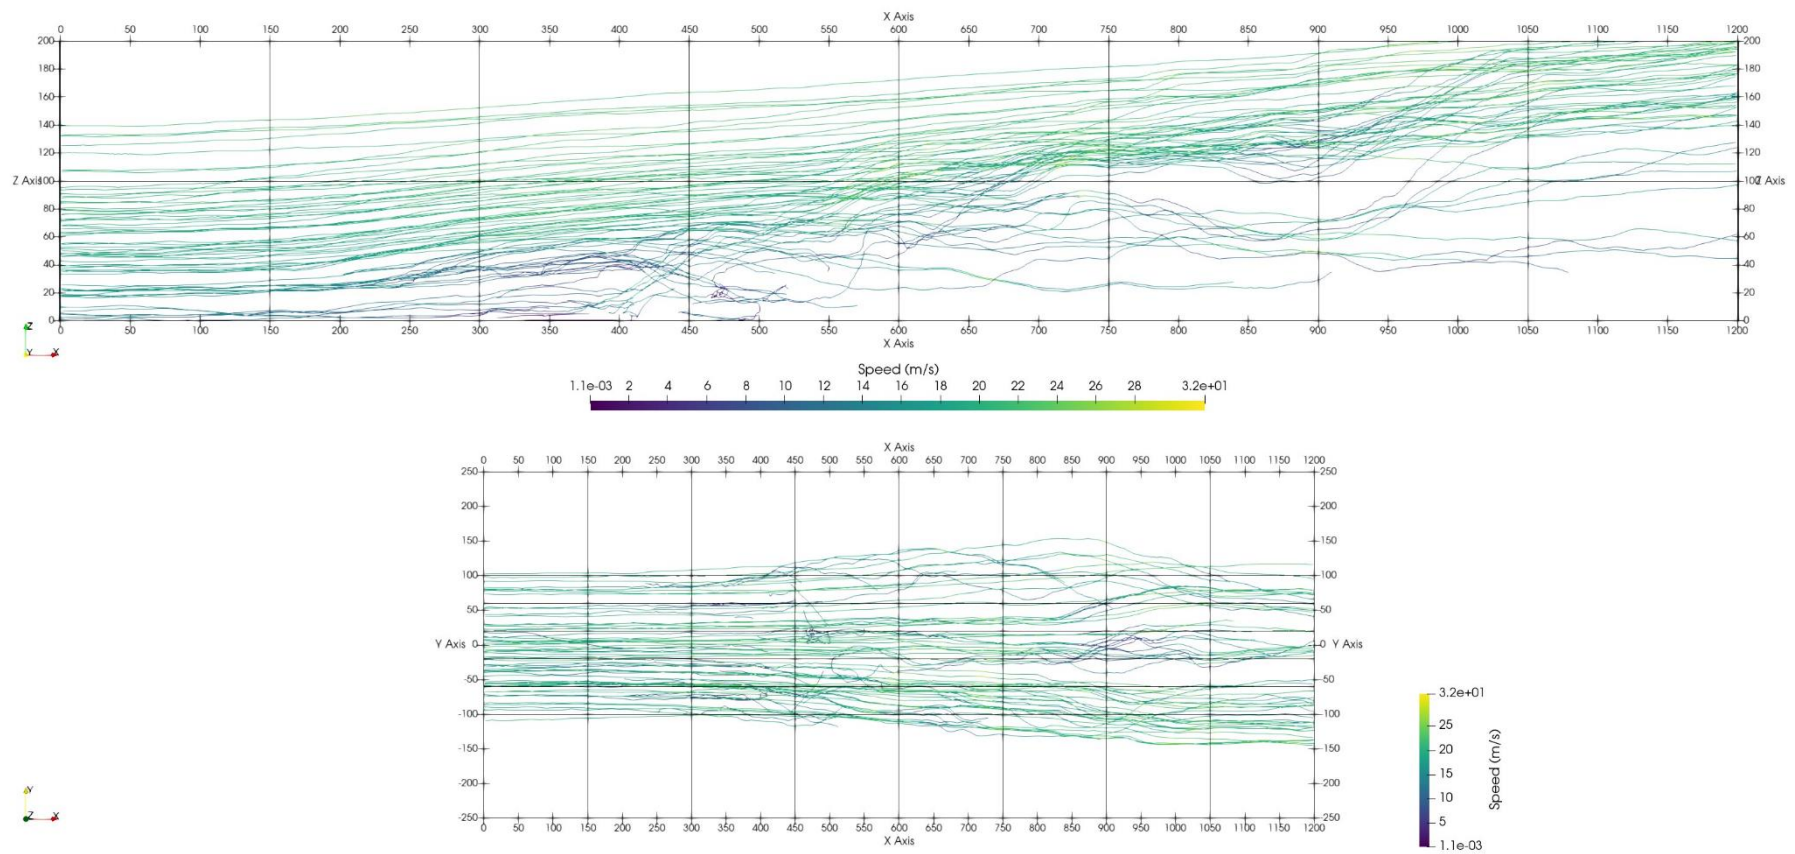


Figure 5-18: Stream tracers for the 54km/hr, 4:1 W:D, Low HRR, Open Field case. Top: View in the X-Z plane. Bottom: View in the X-Y plane.

Comparing each pair of forested and open field cases (e.g. Figure 5-13 and Figure 5-14) we can observe the effect of the forested domain on the path of the stream tracers. In the forested cases, there are generally very few streamlines at low altitude in the region beyond $x = 400\text{m}$ – the area downwind of the fire. In these cases, streamlines that pass close to the burning region either quickly increase in altitude or come to a halt a relatively short distance downwind of the fire. This means that the population of particles distributed over the far downwind regions of the forested domains were either carried forward at low altitudes by momentum gained close to the fire or were lofted to higher altitudes before falling.

Contrasting this, the open field cases show a number of continuous streamlines passing through the fire and exiting at low altitude before continuing through the rest of the domain relatively close to the ground. We can also observe some vortex formation close to the ground in the open field cases, such as those shown in Figure 5-16. While the low-altitude streamlines are moving at lower speeds than higher-altitude streamlines in the open field cases, they have not come to an almost complete halt like in the forested case. This difference in low-altitude streamline paths and velocities is likely to be a significant factor in how the distribution of particles differs between forested and non-forested cases.

Looking at the stream tracers from a top-down perspective across all six cases, we can see that in both the forested and non-forested cases there is a tendency for the streamlines to bulge outwards around the convective plume in the X-Y plane – the stream tracers tend to gradually move away from the $y = 0$ line in the region just downwind of the fire. At lower wind speeds (such as in Figure 5-13 and Figure 5-14) we can see that the streamlines then re-converge towards the end of the domain, producing a high density of stream tracers along the centerline at the $x = 1200\text{m}$ line. At higher wind speeds (as in Figure 5-15 and Figure 5-16) the divergence of streamlines away from the $y = 0$ line is much more extreme, and there is no noticeable re-convergence of streamlines at the end of the domain.

5.5.2 Downwind Dispersion and Initial Height

Having made these observations about the path of massless tracers moving through the domain, we can then look at the characteristics of the path of simulated particles with mass and see if similar trends appear, and how these trends impact the overall distribution of particles. The simplest analysis of particle trajectory is to observe how each parameter impacts the average peak altitude of particles in each scenario and the landing position of particles. However, a large proportion of particles in all scenarios do not leave the area where they were inserted –

some particles fall straight down, or are only carried forward a small distance and land within the area we defined as burning. Since these particles cannot contribute to any further fire spread or potential damage to structures, particles that do not land outside of the burning zone are not included in most further analysis – but we need to account for these particles. A plot of the proportion of particles that escape the burning area for each case is shown in Figure 5-19.




Figure 5-19: Plots of the proportion of particles that leave the burning area for each case.

This shows an interesting difference between the open field and forested scenarios. For the open field scenarios, the proportion of particles that escape the burning area is almost constant

regardless of the heat release rate, wind speed, or W:D ratio. In almost all cases, approximately half of the particles inserted escape the burning area.

Contrasting this, the forested cases show a clear downward trend with the size of the fire. Larger fires, and fires with lower heat release rates, have a consistently smaller proportion of particles that escape the burning area. The effect of fire size seems to level off, with the very largest fires producing a similar proportion of particles leaving the burning area to fires less than half their size. HRR has a fairly consistent effect across all cases – the low HRR cases typically show a reduction of around 27% of the particle fraction leaving the burning area when compared to high HRR cases with otherwise identical conditions.

This difference in behaviour between the open field and forested cases suggests that there is a distinct difference in how particles are transported in these scenarios. In the open field cases, it is possible that many particles that experience relatively little lofting still manage to escape the burning area by being blown directly forward from the point they were inserted due to higher sub-canopy wind speeds – as can be seen in the path of massless tracers for the open field case (e.g. Figure 5-14). Contrasting this, the extremely low sub-canopy wind speeds in the forested cases may prevent this from occurring, meaning that only particles that can be effectively lofted up to a higher altitude will be transported any meaningful distance from the fire. The smaller fires in the forested domain may show a larger proportion of particles escaping the burning area simply due to a larger ratio of perimeter to area – a larger proportion of particles are inserted very close to the edge of the burning area in smaller fire cases.

Taking this difference in particle transport behaviour into account, we can make some observations about the maximum altitude achieved by particles in forested and non-forested cases. The position of each particle that landed outside the burning area was recorded every second from the moment the particle was inserted into the domain, and the location of maximum particle altitude was found. The mean maximum altitude and standard deviation of maximum altitude for open field and forested cases are shown in Figure 5-20.

The most immediately apparent difference in particle maximum altitude between forested and open field cases is the much greater variance seen in forested cases. While the open field cases generally have very consistent maximum particle altitudes, the forested cases have considerably larger variances – sometimes more than triple that of the comparable open field case. In addition to this, many of the forested cases have a slightly higher mean maximum particle altitude than their open field counterparts. This supports the previously discussed

difference in transport mechanism between forested and open field cases – in the forested cases, the higher mean and standard deviation of maximum particle altitude is due to the higher proportion of particles landing outside of the burning area that were lofted by the convective plume. These lofted particles achieve a greater mean altitude than particles that are simply blown straight forward from their insertion point, and the lofting process produces a greater variance than transport in the open field cases.





Figure 5-20: Graph of average maximum particle altitude for particles that left the burning area for forested and open field cases. Blue bar indicates average maximum particle altitude, error bars show 1 standard deviation around the average.

While the presence or absence of forest in the downwind area of the domain makes a very obvious difference in the variance of particle maximum altitude, the impact of the other parameters is much smaller and less clear. Wind speed has a small effect on the average particle

height achieved – a small decrease in average particle height as wind speed increases, likely due to a larger proportion of particles at low altitudes that escape the burning region by being carried forward on the wind with little lofting. Cases with higher heat release rates tended to have a higher maximum particle altitude, but the size of the effect was only ~3.3m averaged across all cases – less than one standard deviation of particle altitude in most cases. Width to depth ratio also had a small negative effect on average particle altitude, with approximately a 3.6m difference between the average particle height achieved between 1:1 and 4:1 W:D cases. Interestingly, cases with a 1:1 width-to-depth ratio had a much larger standard deviation in the x -position at which maximum particle altitude was achieved compared to other cases, especially in the open field scenarios.

The effect of particle initial altitude on the final landing position can be visualized using a combination of data presented in Figure 5-21 and Figure 5-22.




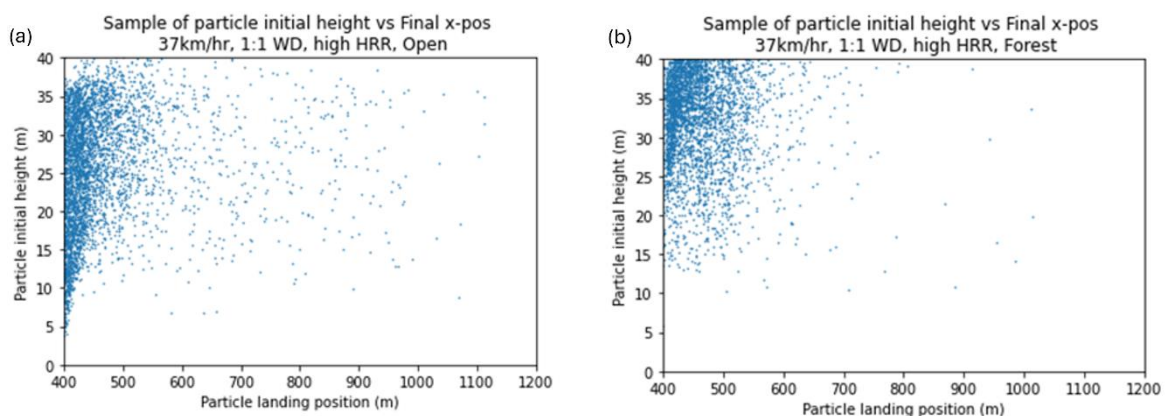


Figure 5-21: Histograms of peak particle altitude, describing what fraction of particles achieve a given maximum altitude under different conditions. Dashed lines indicate treetop altitude.

Figure 5-21 shows a series of histograms of maximum particle altitude achieved under different initial conditions for the set of particles that escaped the burning area. These histograms capture the behaviour of two distinct populations of particles – particles that experience a significant degree of lofting in their trajectory out of the burning region, and particles that are not lofted above their initial position. The particles that experience significant lofting make up the long right-hand tail of the histogram. Particles that do not experience significant lofting make up the bulk of the distribution at 40m and below. The peak of the histogram is consistently around 40m, the height of the top of the canopy and the maximum initial height of any particle. We can determine from the shape of these histograms that particle initial height is one of the most significant factors in determining whether or not a given particle will escape the burning region – the proportion of particles that experience significant lofting is relatively small compared to those that do not. It can also be observed that as the W:D ratio increases, the size of the tail of the histogram tends to decrease – indicating a larger fraction of particles that are not lofted above the canopy (40m). The relationship between particle initial height and particle trajectory can be further investigated by plotting the final landing position of particles in the x direction against initial particle height, as in Figure 5-22.






Figure 5-22: Plots of initial particle altitude against final particle landing position in the x -axis for 37 km/hr cases. Sample of 100 particles for 42 particle types, total of 4200 particles in each plot.

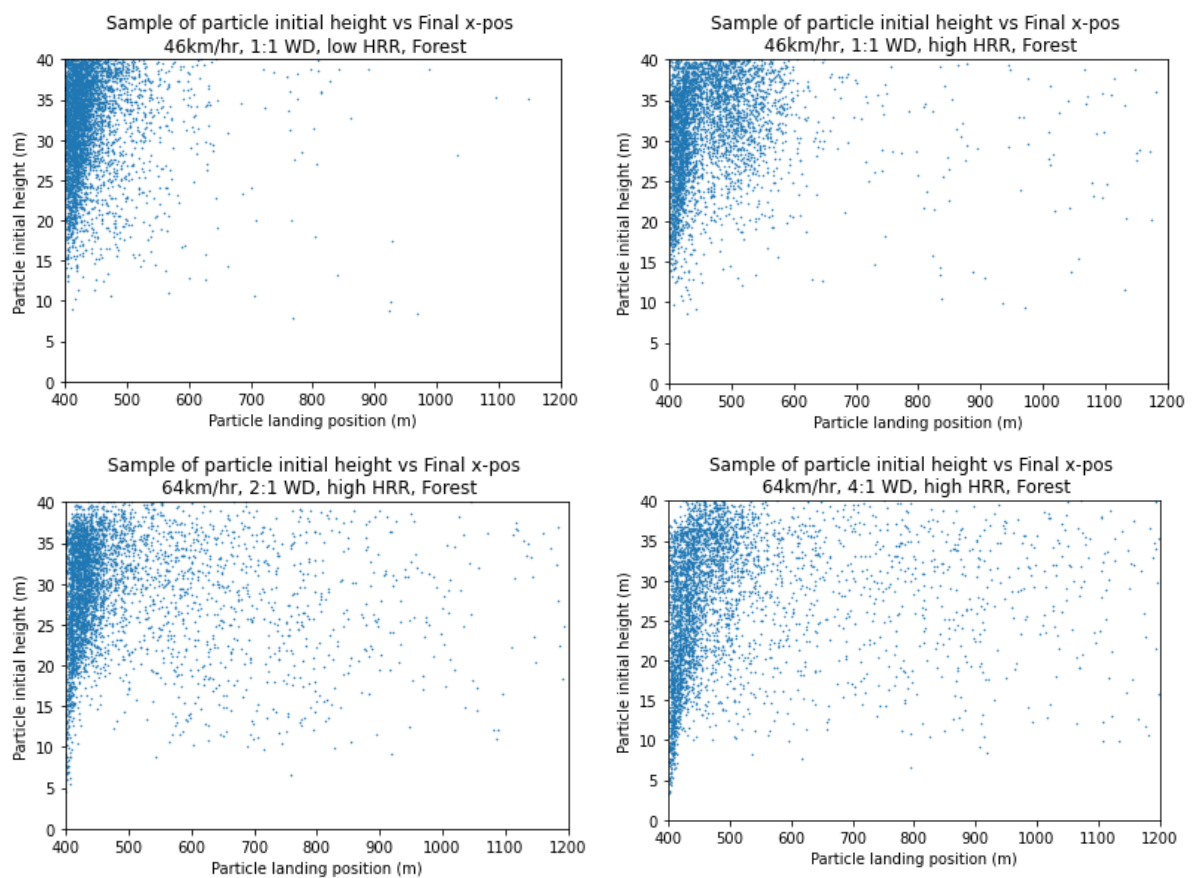


Figure 5-23: Plots of initial particle altitude against final particle landing position in the x -axis for a selection of 46km/hr and 64km/hr cases. Sample of 100 particles for 42 particle types, total of 4200 particles in each plot.

These plots show a clear difference in the relationship between particle initial position and final landing position between forested and non-forested cases, as well as the effect of heat release rate. In the forested cases, the entire distribution is shifted upwards and to the left, with fewer particles dispersed across the length of the x -axis. This indicates that for the particles that did escape the burning region in forested cases, a larger proportion of them started at a higher

altitude, and that those particles were not dispersed as far downwind as their counterparts in open field cases. The effect of heat release rate is also apparent, with the lower HRR cases distributing fewer particles across the far downwind regions.

5.5.3 Crosswind Dispersion

Firebrand particle distribution in the crosswind direction can be investigated using a similar approach to firebrand particle distribution in the downwind direction. A random sample of firebrand particles that escaped the burning region was chosen for each set of conditions, with an equal number of particles chosen from each particle species. These particles' initial and final positions were extracted from the data using a Python script. The difference between the initial and final firebrand particle y -position can be used as a means to observe the dispersion of particles in the crosswind direction that does not simply depend on the width of the fire.

Plots of particle initial and final y -positions can be used to show how the initial position of a firebrand particle impacts its distribution throughout the domain. Examples of these plots are shown in Figure 5-24. In these plots, a solid 45° line would represent a distribution of particles that travel straight forward from their initial position and do not disperse in the crosswind direction at all.

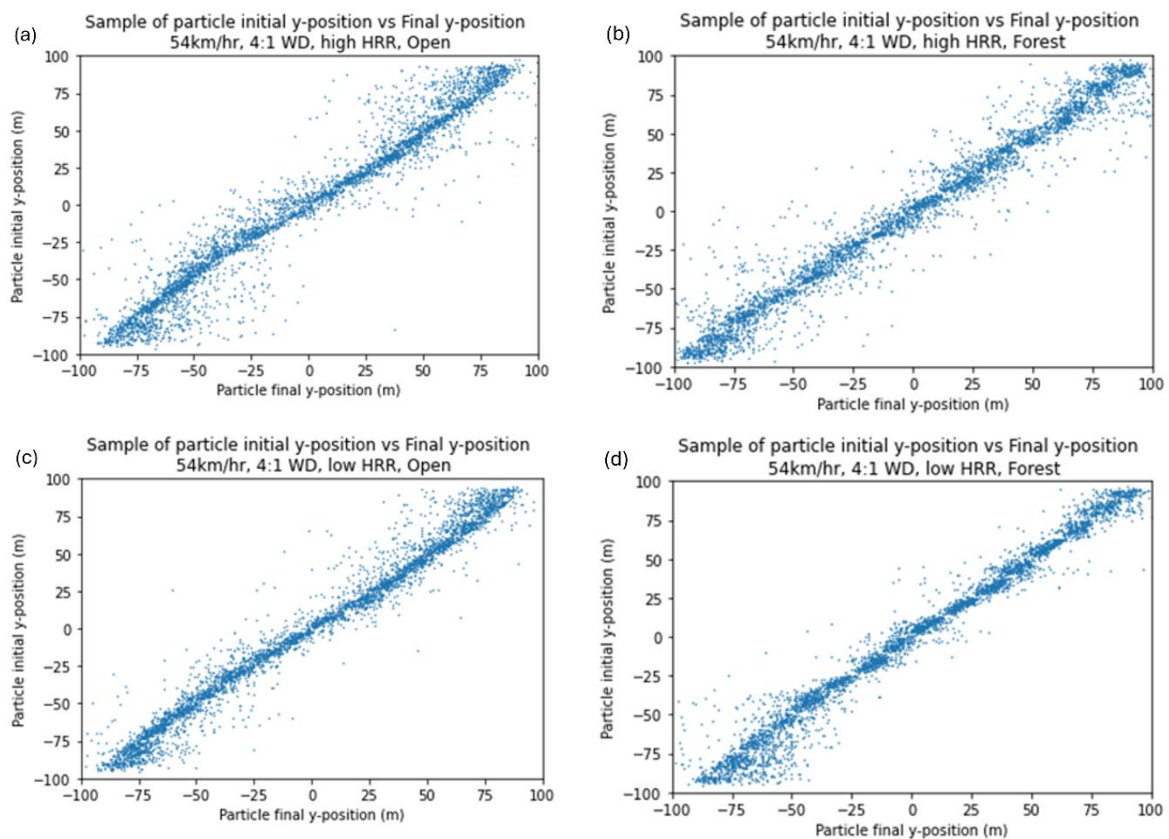
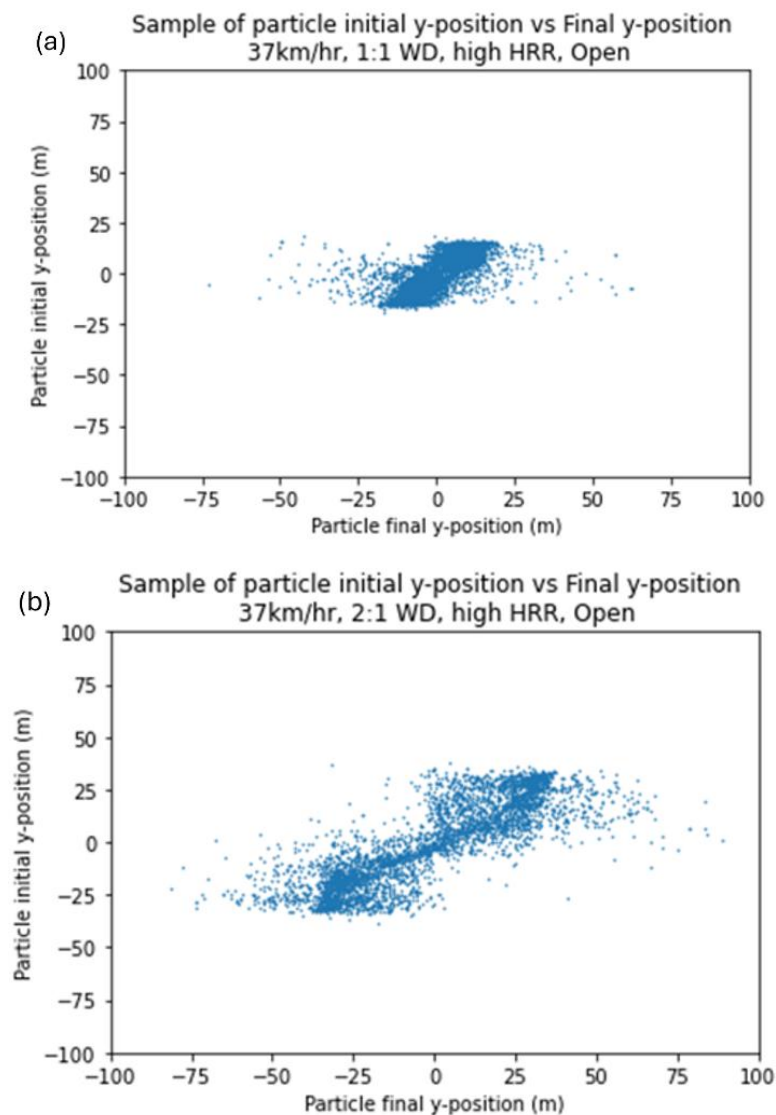


Figure 5-24: Plots of firebrand initial and final positions for a range of 4:1 WD, 54km/hr cases.

Comparing the open field plots of initial and final y -position to comparable forested cases, it can be observed that the open field cases tend to produce a gently curved distribution of particles, whereas the forested cases produce a nearly straight distribution. The shape of the curve in the open field cases reflects the two-peaked distribution observed in the statistical analysis of particle distribution. This effect is most obvious in the wider cases, but can also be observed in some of the narrower cases, as shown in Figure 5-25.






Figure 5-25: Plots of firebrand initial position for 37km/hr, high HRR, open field cases of increasing $W:D$ ratio.

We can more closely investigate this curve of particle distribution by plotting the change in particle y-position for different initial y-positions (Δy), rather than initial and final y-positions, as well as plotting a moving average of Δy over the distribution of particles. In these plots, a straight and horizontal line would represent a distribution of particles that travel straight forward from their initial positions. Comparing the same cases shown in Figure 5-24, we obtain the plots shown in Figure 5-26.

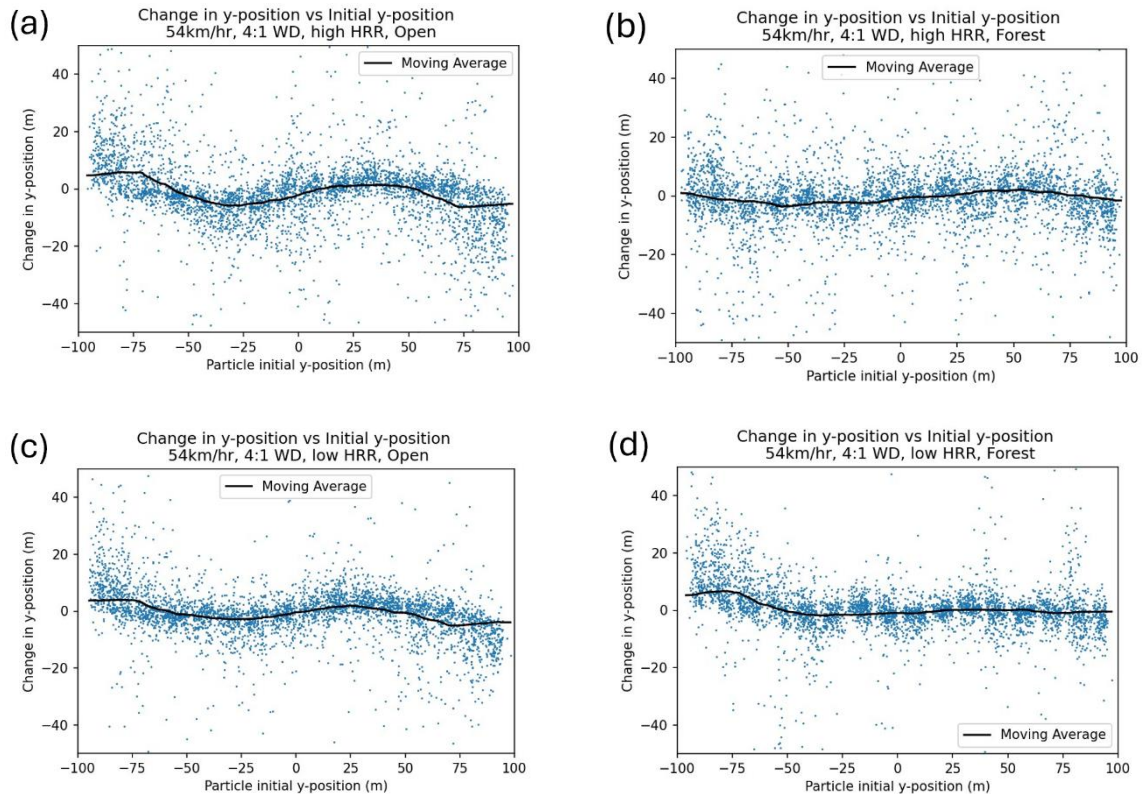


Figure 5-26: Plots of change in particle y-position for different initial positions for a range of 4:1 WD, 54km/hr cases.

By looking at the scatter plot of particle dispersion and the moving average of Δy across initial y positions, we can observe trends in how particles that begin at a given y position tend to move across the domain. In particular, we can see how the two-peaked distribution that is produced in higher W:D open field cases is produced, and how the presence of a forest can impact particle distribution in the crosswind direction to produce single-peaked distributions.

In the open field cases shown in Figure 5-26, it can be observed that the moving average of Δy takes on a roughly symmetrical oscillating shape, with local maximum and minimum Δy values appearing between the extreme edges of the burning region and the center of the domain. A local maximum appears on the positive side of the y axis and the local minimum appears on the negative side, showing a tendency for particles in these regions to move further towards the edges of the domain. The extreme ends of the Δy curve tend to show particles moving inwards, towards the center. These tendencies combine to produce peak particle densities in between the edges of the burning region and the location of the local minima and maxima. At the center of the domain, the average Δy is often close to zero.

To illustrate this trend, consider Figure 5-26a – the 54km/hr, 4:1 WD, High HRR, Open Field case. In this case, we can clearly observe a local minimum Δy in the region around -30m, and a local maximum Δy at a similar position on the positive side of the y axis. This indicates that particles that begin at intermediate positions along the y axis tend to move outwards, away from the center of the domain. At the extreme edges of the distribution, we see a tendency for particles to move inwards, towards the center of the domain – near $\pm 80\text{m}$, we can see the absolute maxima and minima on the Δy curve. These two effects combine to produce the two-peaked shape we observe in this case with maximum particle landing density occurring at roughly $\pm 55\text{m}$, shown in Figure 5-27.




Figure 5-27: Kernel density estimate plot for the 54km/hr, 4:1 W:D, High HRR, Open Field case.

We can then consider the shape of the Δy curve for other scenarios and observe how different initial conditions impact crosswind dispersion to produce the firebrand particle distributions we have obtained. First, we can compare similar open field cases with different HRRs – such as Figure 5-26a and Figure 5-26c. For these two cases, we can see that the lower HRR case (Figure 5-26c) has a similar overall shape to the higher HRR case (Figure 5-26a) but with less extreme values. This is typical of high/low HRR case pairs in open field conditions – usually, the lower HRR case will have a similarly shaped Δy curve, but with less extreme values.

The forested scenarios show an interesting variety of behaviours in the Δy curve. While many forested cases produce nearly Δy flat curves, such as in Figure 5-26b and Figure 5-26d, some forested cases produce Δy curves with more extreme values than comparable open field cases. While these curves indicate that particles in these cases may travel a great distance along the y axis, these scenarios still produce single-peaked distributions, or distributions with relatively little separation between peaks.

An example of this is shown in Figure 5-28b, where the 54km/hr, 1:1 WD, High HRR, Forested case produces a very extreme Δy curve when compared to a similar open field case shown in

Figure 5-28a. The kernel density estimate plot for this case is presented in Figure 5-29 where we can see that the overall crosswind marginal distribution (shown on the right of Figure 5-29) has only an extremely small separation between peaks, rather than a clearly two-peaked distribution. Looking at the scatter plot of particle displacement, we can see how this may occur: many particles in this scenario are displaced extremely far from their initial position, away from the centre of the domain. Rather than settle into two peaks, these particles are distributed across a wide area, widening the overall distribution without producing strong peaks. These particles can be seen in the KDE plot as two outwards-pointing lobes in the region between $x=500$ and $x=600$.

Figure 5-28c and Figure 5-28d show a more typical relationship between open field and forested scenarios. The open field case produces a large, smooth curve of particle displacement and produces a clearly two-peaked distribution in its KDE plot, while the forested case produces a much flatter curve with a large and relatively flat distribution in the KDE plot.

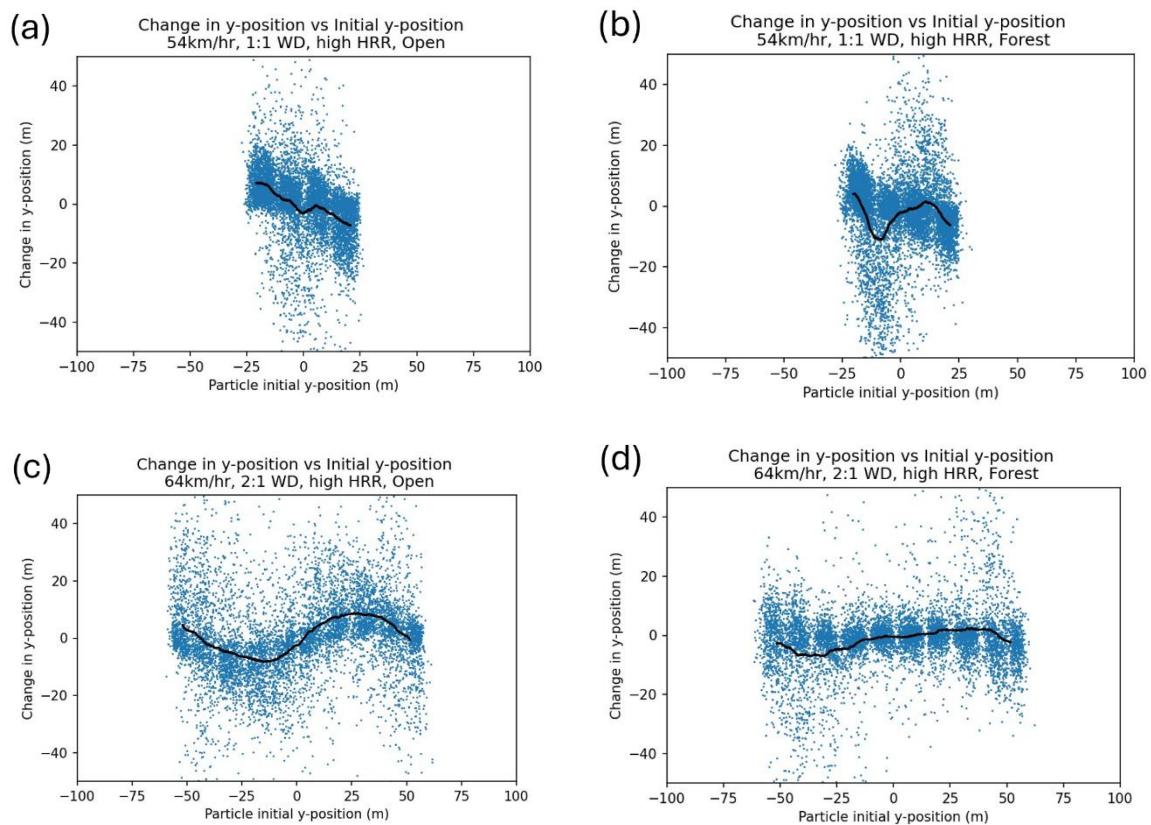


Figure 5-28: Change in particle y-position for a selection of narrower $W:D$ fire cases at 54km/hr and 64km/hr




Figure 5-29: Kernel Density Estimate plot of particle distribution for the 54km/hr, 1:1 WD, High HRR, Forested case. This corresponds to the curve shown in Figure 5-28b.

5.6 Overall particle distribution and trajectory

The difference in overall firebrand particle distribution in each set of scenarios can be explained as a consequence of observed plume, particle, and vortex behaviour. The scenarios can be divided into three categories based on the forest condition of the domain and the observed shape of the distribution:

1. Open Field scenarios where the crosswind marginal distribution is two-peaked
2. Forested scenarios where the crosswind marginal distribution is single-peaked or flat
3. Open Field scenarios where the crosswind marginal distribution is single-peaked

In the first category of scenarios (open field scenarios with a two-peaked distribution), the convective plume is bent over at a steeper angle than predicted by the Zonato model of plume rise – as shown in Figure 5-2a. As the counter-rotating vortex pair is formed beneath the plume, this steeply bent-over plume produces a pair of counter-rotating vortices that sit at a lower altitude than other cases, especially in the region very close to the fire (i.e. less than 50m downwind). In addition to this, the distribution of maximum particle altitude achieved in open field cases peaks at 40m (see Figure 5-21) with only a small proportion of particles that achieve an altitude of over 40m. This indicates that most particles are only lofted to altitudes close to

the top of the canopy, or are not lofted at all, decreasing in altitude from the moment they are inserted. These factors interact to produce the observed Δy curves in these cases, and the resulting two-peaked distribution of particles. With the bulk of particles escaping the burning area at altitudes between 20-40m, these particles are very likely to pass through the counter-rotating vortex pair in the area immediately downwind of the fire.

When a particle passes through the CVP at these altitudes, it follows a particular trajectory that can be identified by plotting its path on the X-Y plane. The particle will at first move outwards, as it passes through the upper region of the CVP where the v -velocity component of the wind points towards the outside edges of the domain. As the particle falls, the v -velocity component of the wind around the particle decreases in magnitude before changing direction, pointing towards the center of the domain. The particle then changes direction, moving inwards until it touches the ground. This can be seen in Figure 5-30, where the trajectory of a random sample of particles is plotted from their point of origin (red) to final landing position (blue) for a range of two-peaked open field cases.







Figure 5-30: Plot of particle trajectories on the X-Y plane for a selection of two-peaked distributions. Initial particle positions are shown in red, final landing positions are shown in blue. Grey lines show the path taken by each particle.

In contrast to open field cases, forested cases tend to have a significantly steeper plume, and we observe that the CVP in forested cases tends to form at higher altitudes than in open field cases. The presence of the forest at low altitudes also interferes with the formation of the vortex pair near the ground. While we observe more lofting in these forested cases with a steeper and faster convective plume, most particles still experience relatively little lofting. This means that, unlike the open field cases, many of the particles in the forested cases that are blown forwards with little lofting may not be at the right altitude to interact with the upper regions of the vortex pair. Instead, they are blown straight forward with randomly directed dispersion, or interact with the lower region of the vortex pair and move inwards. A fraction of particles that are carried upwards in the convective plume reach high enough altitudes to interact with the CVP, and these particles are often displaced a great distance outwards towards the edges of the domain. While these particles that enter the CVP in forested scenarios show a similar trajectory to particles in open field cases, the lower v -velocity components beneath the canopy do not allow these particles to travel as far inward at low altitudes. This results in a wider dispersion of particles and less well-defined peaks. This can be seen by comparing trajectory plots for forested cases against similar open field cases, as in Figure 5-31.

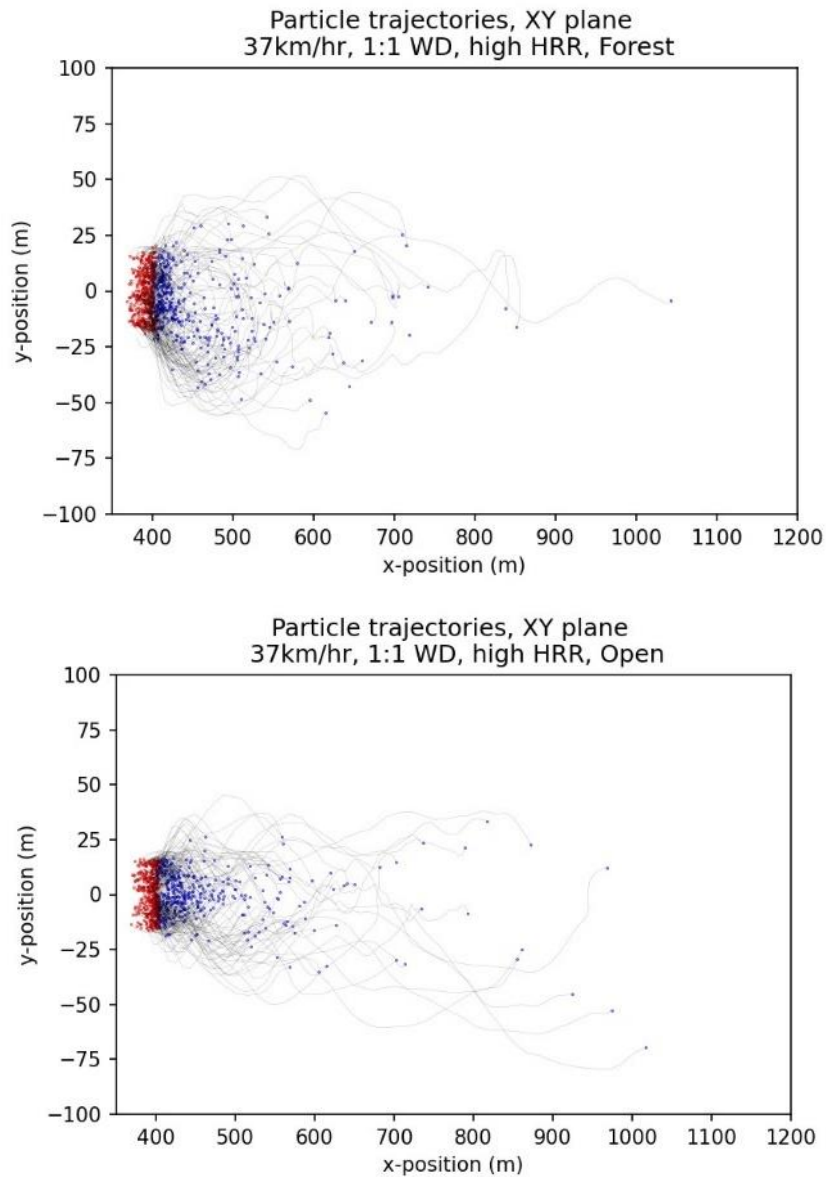


Figure 5-31: Comparison of particle trajectories between similar forested and open field scenarios. The forested case produces a wider and more even distribution of particles. Initial particle positions are shown in red, final landing positions are shown in blue. Grey lines show the path taken by each particle.

Further examples of particle trajectories in forested cases are shown in Figure 5-32:







Figure 5-32: Plot of particle trajectories on the X-Y plane for a selection of forested cases, showing a wider and more uniform distribution of particles. Initial particle positions are shown in red, final landing positions are shown in blue. Grey lines show the path taken by each particle.

The third category of scenarios – open field scenarios that produce single-peaked distributions – can be understood easily by studying particle trajectories. Particles in these scenarios follow a similar trajectory to particles in other open field scenarios. The characteristic outward then inward motion of particles moving through the CVP is present, and the altitude of the CVP in these cases is similar to that of the two-peaked cases. However, while the particles follow a similar trajectory, the narrow width of the fire in these cases leads to a significant overlap in particle landing position regardless of which side of the center line a particle originates from. The single peaked distributions obtained in these cases are effectively two peaks that overlap so greatly that they are difficult to distinguish. This can be seen most clearly in low HRR cases, as shown in Figure 5-33.

Collectively, these observations about the impact of changes in plume and vortex characteristics on the trajectory of a population of firebrands across the range of scenarios may constitute a novel observation. This is due to both the specific bimodal distribution that we have observed and the identification of the physical processes that are likely to produce this distribution.





Figure 5-33: Plot of particle trajectories on the X-Y plane for a selection of single-peaked open field cases, showing a large degree of overlap in final landing position for particles that originate on opposite sides of the center line. Initial particle positions are shown in red, final landing positions are shown in blue. Grey lines show the path taken by each particle

Chapter 6 - Statistical Methodology and Model Development

Contents

Introduction.....	141
Data Visualization and Qualitative trends	142
Histograms	142
Kernel Density Estimation.....	140
Statistical Trends in Particle Distribution	141
Downwind Axis	141
Crosswind Axis.....	144
Trends across particle dimensions and geometry	151
Curve Fitting	150
Downwind Marginal Distribution.....	156
Crosswind Marginal Distribution	162
Describing the Joint Distribution.....	162
Copula Method.....	162
Conditional Probability Method	164
Model Fitting	171

6.1 Introduction

In order to develop the simplified model of firebrand landing, it is necessary to investigate the statistical properties of the distribution of particles landing over the domain in each scenario. By describing the distribution of particles statistically, the impact of the different initial physical conditions on particle landing position can be described as a change in the parameters of a set of probability density functions.

The data used for this statistical analysis is primarily data relating to particle position. For each of the particles simulated in each scenario, particle positions were recorded once per second of simulated time. Each particle is uniquely identified with an eight-digit ID number and is classified by particle geometry and dimensions into 42 distinct particle species. This produces a collection of output files that record the trajectory of every individual particle from the moment they are introduced into the simulation domain in one-second increments. These particle output files are generated by FDS in a binary format, then converted into a human-readable CSV format using a Python program.

The position of particles on the ground at the end of the simulated time (180 seconds) is assumed to be the final landing position of particles, with any particles that may still be airborne at that time excluded from analysis. This is done by filtering particle position data at the final time step and removing all particles with vertical positions more than $\frac{1}{2}$ of a grid cell above the ground. This filtered set of particle positions at the final time step is used as the basis for the majority of subsequent analysis.

6.2 Data Visualization and Qualitative trends

6.2.1 Histograms

To obtain a broad understanding of the shape of the particle landing distribution, it is necessary to visualize the data. While the simplest visualization of the data may be a simple scatter plot, the large number of particles would make such a visualization a poor choice. Instead, a 2-D histogram with a fixed number of bins over the area downwind of the fire was used. Applying this visualization to the filtered dataset at the final time step and using a logarithmic colour scale for density, we can obtain a series of histograms like those shown in Figure 6-1.




Figure 6-1: Histogram of firebrand landing for 46km/hr, High HRR, 1:1 Width to depth, open field scenario. The colour gradient shows the log of density.

While these histograms are a very basic form of data visualization, they are a useful first step towards a qualitative understanding of trends in particle distribution. Using these histograms, we can visually compare particle landing distributions between scenarios and observe some general trends. For example, we can observe a clear change in the shape of the distribution as the W:D ratio increases:

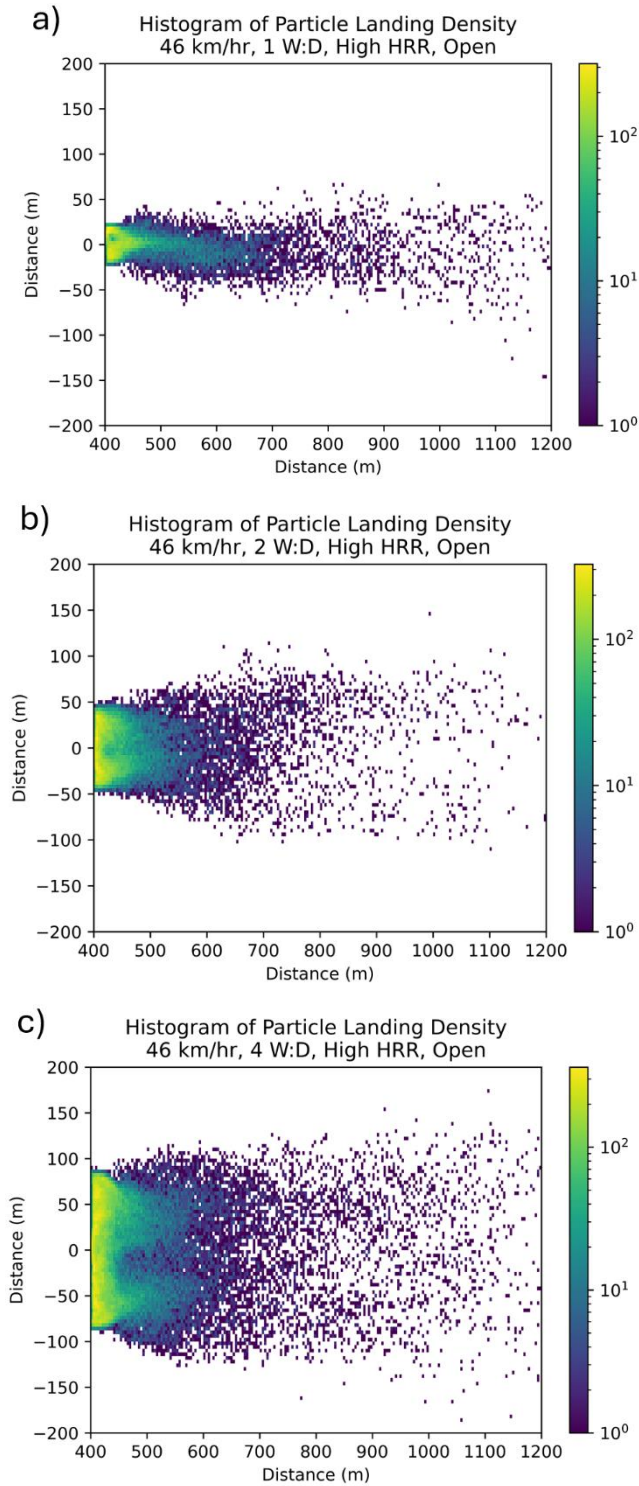


Figure 6-2: Histograms of firebrand landing for 46km/hr, High HRR, open field scenarios. a) 1:1 W:D ratio, b) 2:1 W:D ratio, c) 4:1 W:D ratio

In this case, we can observe that as the W:D ratio increases, the maximum landing density stays almost constant, the distribution widens, and an area of lower particle landing density is observed along the center of the distribution. We may also observe how the shape of the distribution changes with HRR:

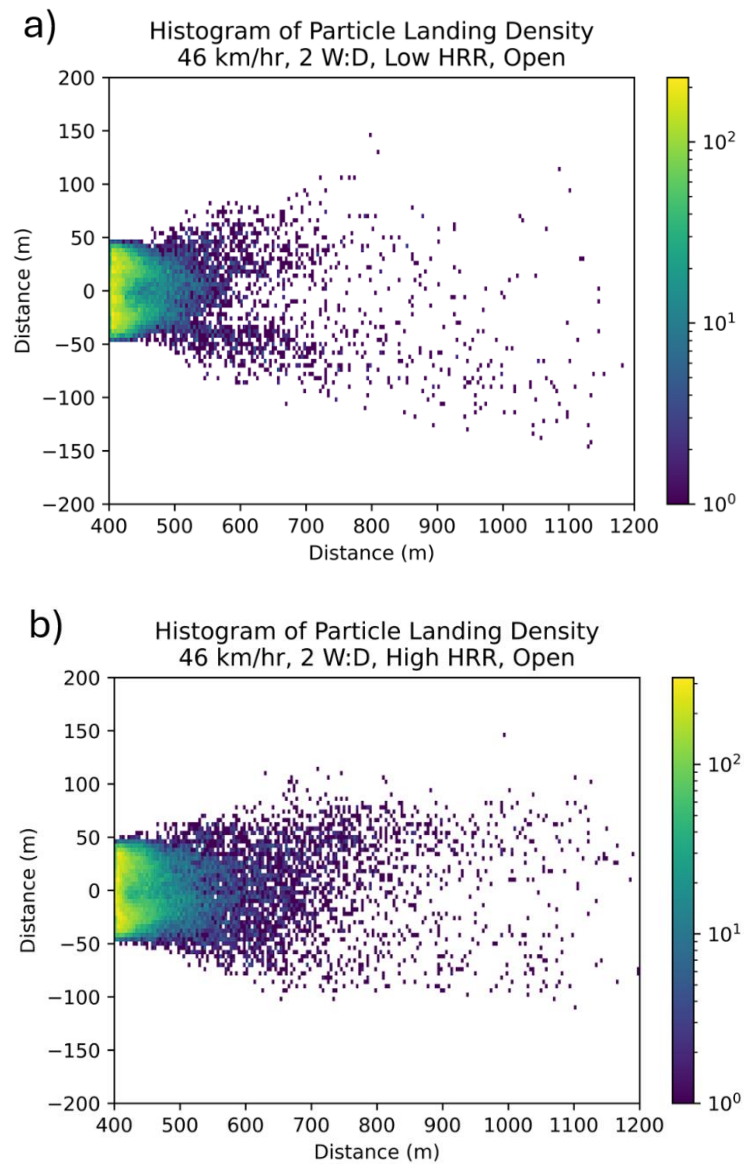


Figure 6-3: Histograms of firebrand landing for 46km/hr, 2:1 W:D ratio, open field scenarios. a) Low HRR, b) High HRR

We can observe that the low HRR case has a similarly shaped distribution to the high HRR case within the first 150m of the fire front, but with a lower peak firebrand landing density. The low HRR case also appears to have a slightly wider distribution in the region beyond 150m. We can also observe that the high HRR case projects a considerably larger number of firebrands further downwind of the fire front.

Although some trends are easy to observe in these histogram visualizations, they are of limited utility if we want to observe subtler trends in our data – especially in the more sparsely populated regions of the distribution, where the histogram is little more useful than a scatter

plot. While they are simple and easy to understand, more advanced visualization techniques can provide additional insight into the shape of our distribution.

6.3 Kernel Density Estimation

Kernel Density Estimation (KDE) is a statistical tool for approximating the probability density function of a random variable [90]. Given a collection of n samples and a kernel function, kernel density estimation approximates the underlying probability density function as a sum of n separate kernel functions. This process produces a plot that somewhat resembles a histogram, but with a smooth gradient of probability density rather than an integer number of samples per bin. For our purposes, a gaussian kernel function was used. The purpose of choosing a gaussian kernel function is to produce a smooth gradient of probability density in regions where data points are sparse.




Figure 6-4: An example of a Kernel Density Estimate plot with marginal distributions. Data is from the open field 46km/hr, High HRR, 4:1 WD case

Because the probability density function approximated by this process is continuous over the region specified, it is convenient to draw probability density contours using this plot. As shown in Figure 6-4, the probability density contours are densely packed in the region nearest to the fire, before gradually spreading out. The final probability density contour covers a much larger area than all the previous contours put together. This final contour includes all particles, including the most widely dispersed outliers – producing a very large contour of very low average density.

Kernel Density Estimation also provides a convenient and smooth approximation of the marginal distributions of the particle landing distribution. The marginal distributions of the KDE plot (top and right of Figure 6-4) are an approximation of the probability density function of particle landing in each axis.

Qualitatively, the marginal distribution in the downwind direction shown in Figure 6-4 is a long-tailed distribution, with a rapid increase in density close to the fire followed by a more gradual decrease. The marginal distribution in the crosswind direction shows a bimodal distribution of probability density, with well-defined peaks separated by a region of low density. The probability density in the crosswind direction has a very short tail, with near-zero density at the extremes. This distribution is symmetrical around the centre line of the domain.

6.4 Statistical Trends in Particle Distribution

In order to develop an adequate model of particle distribution, we must take a quantitative approach to describe the effect of different physical parameters on the distribution of particles. From earlier qualitative analysis we know that the shapes of the marginal distributions are very dissimilar, and that the shape of the overall distribution changes significantly between cases – primarily dependent on the width to depth ratio of the fire. Taking this into account, the following analysis will focus on describing some of the statistical properties of particle positions in each axis separately before moving towards a description of the distribution as a whole.

6.4.1 Downwind Axis

Particle distribution in the downwind direction is affected by all parameters to varying degrees. The effect of changes in each parameter on average particle downwind travel distance can be seen in Figure 6-5:

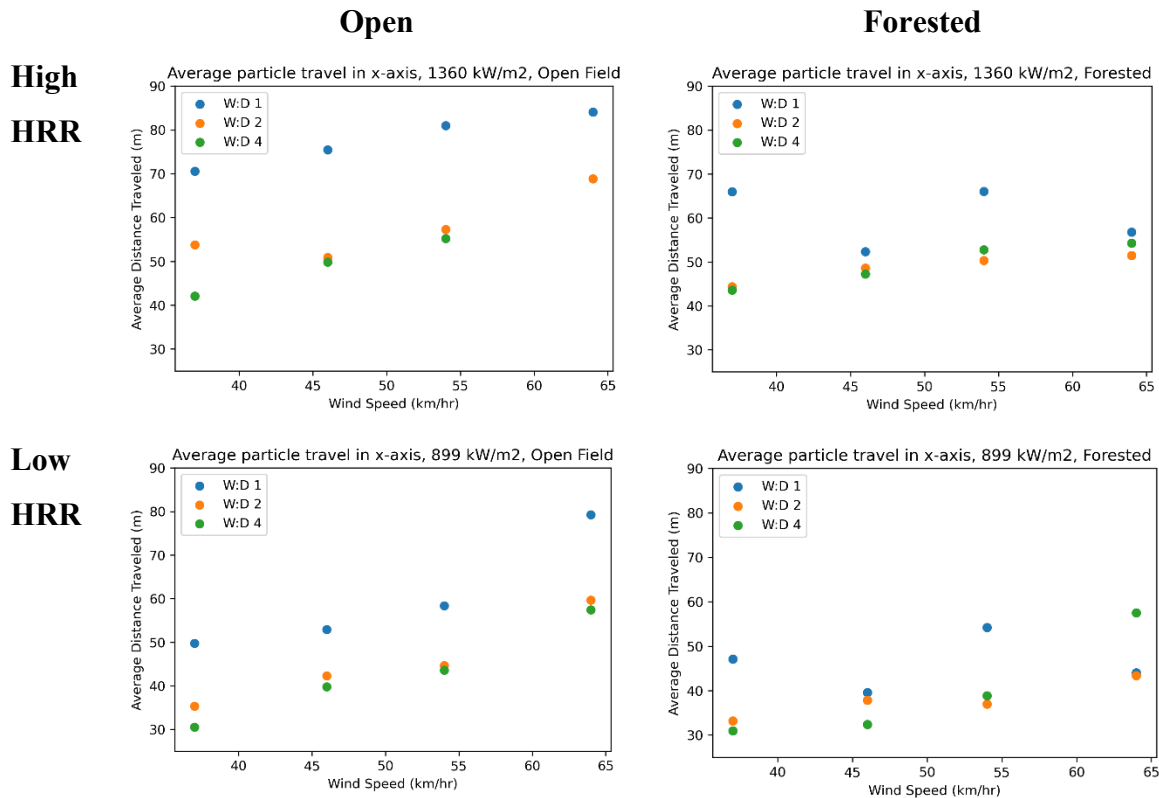


Figure 6-5: Trends in average downwind distance travelled.

As expected, increases in wind speed tend to increase the average distance travelled downwind. Interestingly, the effect of increased wind speed is quite limited relative to the impact of changes in W:D ratio. Scenarios with a 1:1 W:D ratio in open terrain consistently produce average particle travel distances significantly larger than 2:1 or 4:1 cases of the same wind speed. The effect of 1:1 W:D on particle travel distance in forested terrain is smaller, but still notable. The 1:1 W:D cases in the forested domain also show a peculiar trend in average travel distance as wind speed increases – rather than showing a clear increasing trend like the majority of other cases, some cases produce larger than expected average travel distances, resulting in an inconsistent trend across wind speeds for this subset of cases.

The effect of increased heat release rate on downwind particle travel distance is large, and can be observed in both the average of particle travel distance, and the extremes of the distribution – as measured by the 10th decile of particle travel distance. The 10th decile distance is impacted more significantly by the increase in heat release rate than the average. This is likely due to the stronger convective plume in higher HRR cases transporting a larger proportion of heavy particle species out of the flaming region and into the distribution, increasing the number of particles landing relatively close to the fire. Simultaneously, the higher HRR also results in increased lofting of lighter particles to higher altitudes, allowing them to be transported further

downwind. This effectively “flattens” the distribution somewhat, extending the tail of the distribution further downwind while changing the average by a smaller amount.

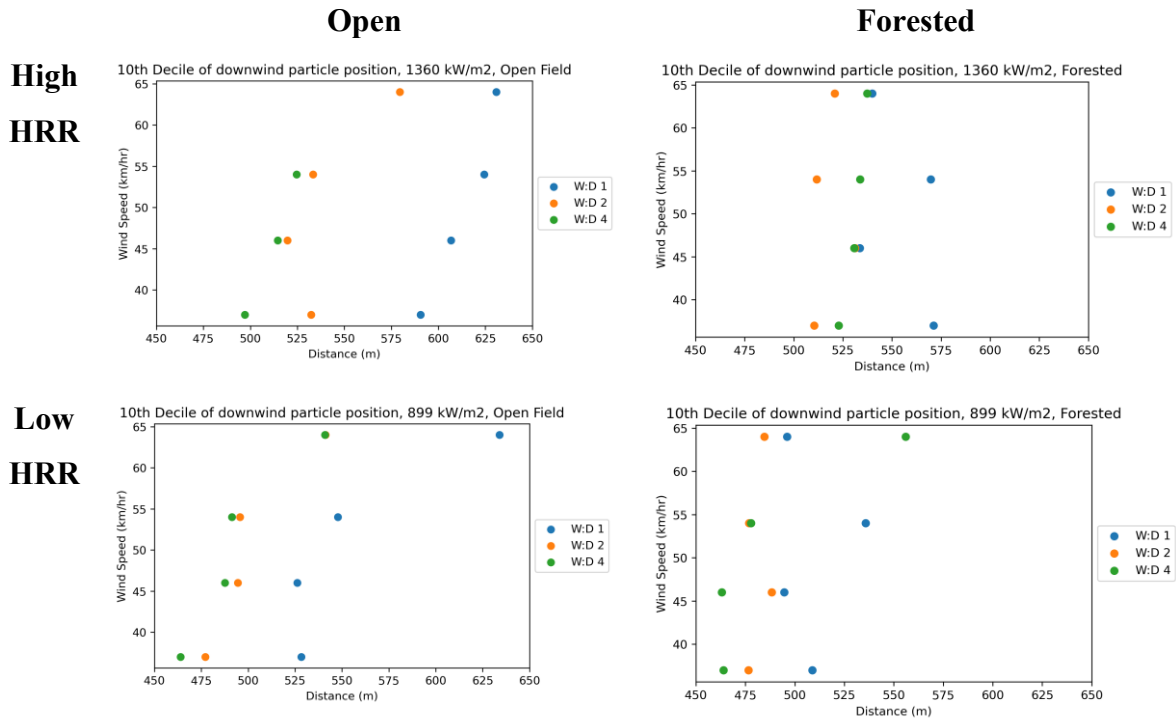


Figure 6-6: Trends in the 10th decile of downwind distance travelled.

The effect of the forest on downwind particle distribution is much harder to capture in simple statistics. In general, it can be observed that the forested cases have a slightly lower average downwind transport distance, and that the cases are grouped more closely together. The presence of the forest also reduces the strength of the effect of wind speed and W:D ratio on the average and 10th decile of downwind transport distance.

We can also examine how the physical parameters affect the variance and coefficient of variation of particle transport in the downwind direction. The coefficient of variation (CV) is a statistic that describes the ratio of the standard deviation to the average of the data – providing a dimensionless number that can be used to compare the variability of distributions between cases with different averages. A plot of the CV statistic is presented in Figure 6-7:

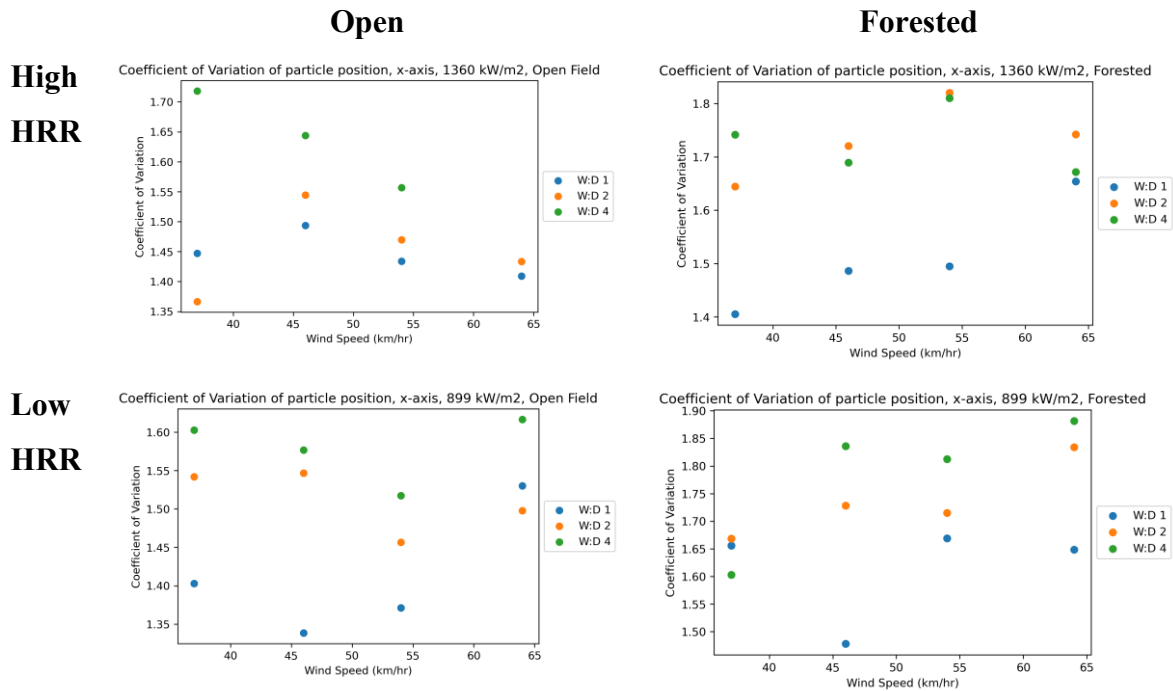


Figure 6-7: Plot of the coefficient of variation in the x-axis for different scenarios.

It can be seen in Figure 6-7 that cases with a W:D of 1 generally have a lower CV statistic than other cases. This is despite the W:D = 1 cases having a greater average particle travel distance and greater 10th decile average. It seems that increasing W:D not only decreases average particle travel distance, but also increases the randomness of the distribution.

6.4.2 Crosswind Axis

The distribution of particles in the crosswind direction is considerably more complex than in the downwind direction. In general, the shape of the marginal distribution in the downwind direction is consistent – this is not the case for the marginal distribution in the crosswind axis. Due to the width of the burning region being a parameter (W:D ratio and wind speed both impact the real width of the fire) examining the average distance from the centreline or variance of the distribution is not useful – the trend is dominated by the effect of the increased width of the fire, as shown in Figure 6-8:

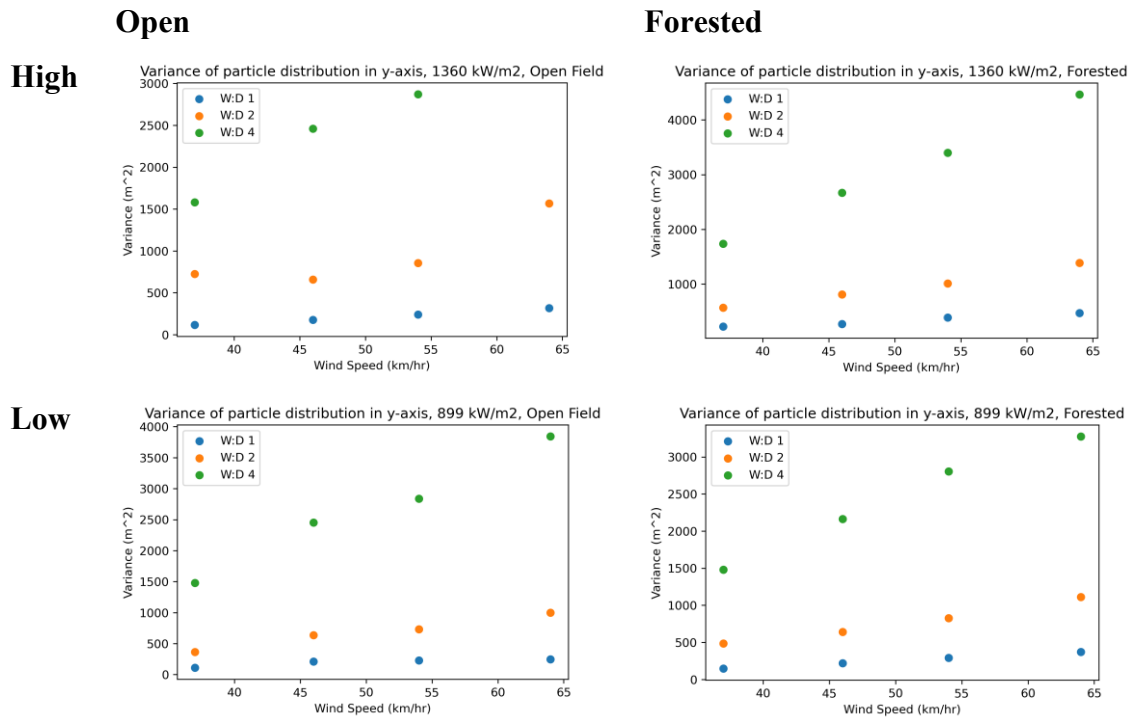


Figure 6-8: Variance of particle distribution in the y-axis, showing linear trends with flame width

As such, we need to use more complex tools to analyse the relationship between the distribution in the crosswind axis and the physical parameters of each scenario. A good starting point is the dimensionless Coefficient of Variation, mentioned previously. To calculate the CV statistic, the average of the absolute value of position in the crosswind axis was used – measuring distance from the centreline of the domain.

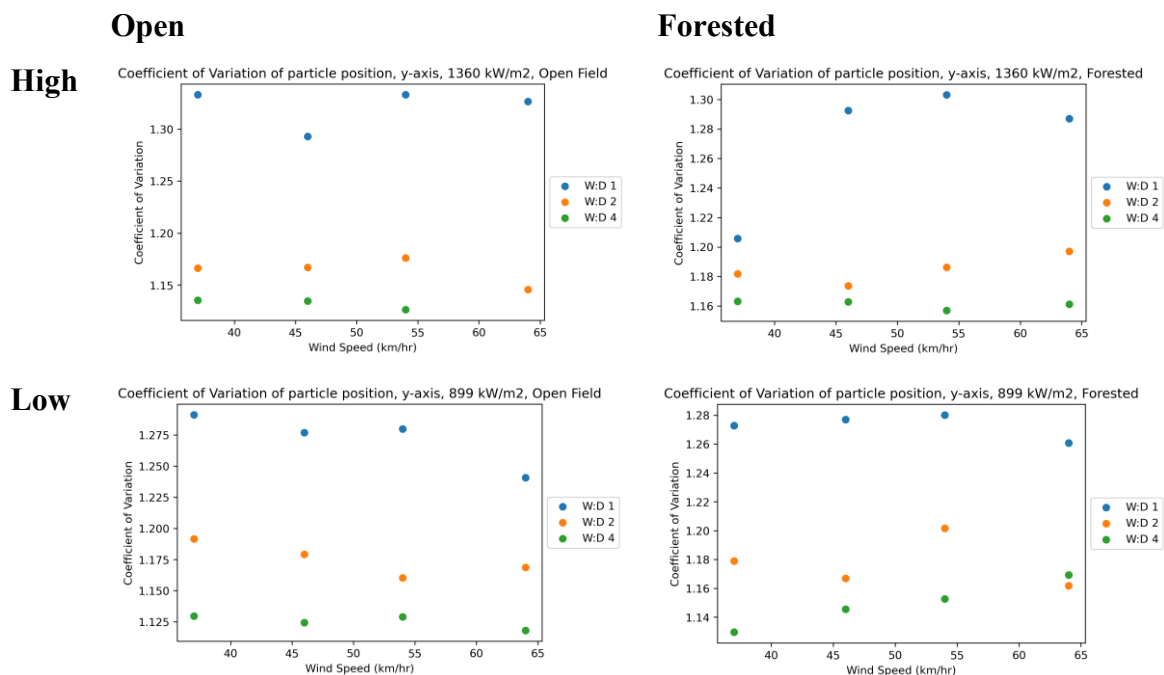


Figure 6-9: Plot of the Coefficient of Variation in the y-axis

Observing the trends in this statistic, we can see that the effect of WD ratio is the dominant source of variation in particle landing position. The coefficient of variation is very similar across wind speeds, HRRs, and even forest conditions, but shows a consistent and clear trend with WD ratio. Despite the larger WD cases having a dramatically larger variance statistic, the narrower WD cases have more actual randomness in particle distribution relative to the average distance travelled in this axis.

6.4.3 Trends across particle dimensions and geometry

Having characterized the whole population of particles in both the downwind and crosswind axes, we can begin to investigate the distributions of individual particle species. As described in the scenario development chapter, there are 42 particle species. These species are divided into three categories – cylindrical, spherical, and rectangular prism (referred to as “cartesian” in FDS). The dimensions and masses of these particle species are based on a selection of experimentally obtained firebrands from Filkov et. al. [91]. The exact dimensions and masses of these particles are presented in Table 6-1, Table 6-2, and Table 6-3 below:

Particle ID	Geometry Type	Radius (mm)	Length (mm)	Mass (g)
2	Cylindrical	1.180	8.770	0.016
3	Cylindrical	1.410	10.1000	0.026
4	Cylindrical	1.330	9.330	0.022
5	Cylindrical	1.045	10.440	0.015
6	Cylindrical	1.420	10.590	0.028
7	Cylindrical	7.060	61.080	4.017
8	Cylindrical	4.250	31.660	0.755
9	Cylindrical	4.085	33.330	0.734
10	Cylindrical	4.610	34.260	0.961
11	Cylindrical	2.965	22.970	0.266
12	Cylindrical	3.115	25.570	0.327
13	Cylindrical	2.915	34.110	0.382
14	Cylindrical	2.645	22.300	0.206
15	Cylindrical	2.040	18.190	0.100
16	Cylindrical	1.425	26.890	0.072
17	Cylindrical	6.110	91.300	4.497
18	Cylindrical	1.450	62.500	0.173
19	Cylindrical	0.700	47.800	0.031
20	Cylindrical	1.515	33.430	0.101
21	Cylindrical	1.900	64.440	0.307
22	Cylindrical	0.625	33.950	0.017
23	Cylindrical	2.175	33.500	0.209

Table 6-1: Dimensions and masses for cylindrical particles.

Particle ID	Geometry Type	Radius (mm)	Mass (g)
24	Spherical	3.510	0.076
25	Spherical	2.890	0.042
26	Spherical	15.660	6.756
27	Spherical	8.610	1.123
28	Spherical	7.640	0.785
29	Spherical	6.670	0.522
30	Spherical	5.280	0.259

Table 6-2: Dimensions and masses for spherical particles.

Particle ID	Geometry Type	Length (mm)	Width (mm)	Depth (mm)	Mass (g)
31	Rectangular Prism	5.740	8.080	3.000	0.059
32	Rectangular Prism	5.450	8.520	3.000	0.059
33	Rectangular Prism	5.430	8.260	3.000	0.057
34	Rectangular Prism	6.580	9.620	3.000	0.080
35	Rectangular Prism	23.260	36.710	3.000	1.084
36	Rectangular Prism	14.600	20.120	3.000	0.373
37	Rectangular Prism	13.840	19.150	3.000	0.336
38	Rectangular Prism	14.650	23.070	3.000	0.429
39	Rectangular Prism	9.750	15.200	3.000	0.188
40	Rectangular Prism	9.680	14.650	3.000	0.180
41	Rectangular Prism	9.190	13.840	3.000	0.161
42	Rectangular Prism	8.510	14.030	3.000	0.152
43	Rectangular Prism	8.610	11.670	3.000	0.128

Table 6-3: Dimensions and masses for rectangular prism particles.

By generating a box plot of the downwind particle travel distance, we can observe several interesting trends in downwind particle distribution. Each particle geometry class (cylindrical, spherical, and rectangular prism) shows a different trend, with particles of similar mass but different geometry exhibiting dramatically different distributions in the downwind axis.




Figure 6-10: Box plot of particle distribution in the downwind direction, separated by particle species. This plot shows the distribution of each particle type in the downwind direction – showing that certain species and geometries are more prone to producing long-tailed distributions than others. In particular, lightweight spherical and cylindrical species produce much longer tails than other species.

Particle species 2 to 23 are cylindrical, with a range of masses from 0.015g to 4.49g. This category includes some of the very lightest particles in the scenario, such as species 2-6 and 19, 20, and 22. These particles also have some of the lowest mass for their cross-sectional area of any particle. Within this category there is very little difference in the median particle travel distance, even between particles of considerably different mass and cross-sectional area. For example, particles 8 to 16 exhibit a range of masses from 0.75g to 0.072g – an entire order of magnitude – but the median travel distance of species 16 is only 10m greater than species 8 (422m and 412m respectively). Even one of the lightest cylindrical species, species 22, has a median travel distance of only 437m. Where particle mass and cross-sectional area seem to make the greatest difference for cylindrical particles is in the tail of the distribution. As can be seen in Figure 6-10, the tails of the distributions consistently grow longer as particle mass decreases.

Particle species 24 to 30 are spherical. These particles exhibit drastically different behaviour to other particle geometries, making up much of the overall distribution in the regions further downwind of the fire. Particle species 24 and 25 – the lightest spherical particles – have a vastly greater average travel distance and a greater inter-quartile range than any other particle species and are the most significant contributor of particles landing at extreme distances. The median travel distance of the spherical particle species shows a stronger correlation with particle mass than other geometries, showing a clear trend of increasing median travel distance as the particles decrease in mass. All of the spherical species exhibit a long-tailed distribution, including the heaviest spherical species. Despite having a greater mass than any other species, particle species 26 has a considerably longer tail than many cylindrical particles with less than half its mass.

Particle species 31 to 43 are rectangular prisms. These particles are notable for having extremely consistent distributions despite changes in geometry and mass. Despite ranging from 1.08g to 0.08g, these particles all have nearly identical median travel distances and tail distributions. This is likely due to the uniform thickness of these particles, representing small sections of bark.

These trends in downwind travel distance across particle species described here appear to hold for all combinations of initial conditions.




Figure 6-11: Box plot of particle distribution in the crosswind direction, separated by particle species. The long-tailed distribution of cylindrical particles can be seen in the plots of particle IDs 24 through 30.

Having described the trends across species for travel in the downwind axis, we may now consider how particle geometries influence transport in the crosswind axis, as shown in Figure 6-11. The distribution of particles in the crosswind direction is remarkably uniform between particle species. Despite the dramatic differences in the shape of the distribution in the downwind direction, the only notable trend in the crosswind distribution is that the spherical particles once again have a larger tail and more outliers in the distribution than other species. However, the difference is not nearly as large as in the downwind direction, and the inter-quartile range of spherical particle species is not noticeably larger than other species.

6.5 Curve Fitting

In the previous section, we described how certain statistical properties of the particle distribution change depending on the physical conditions set at the start of the simulation, and how these statistical properties vary between different particle geometries and masses. Taking these into account, we can begin to develop a statistical model of the distribution as a whole. From the earlier analysis, we can see that the marginal distributions in each axis are very distinct – as such, the process of fitting a function to describe the shape of each marginal

distribution will be considered separately. We will then investigate the development of a coupling function to combine the marginals and produce a preliminary model.

6.5.1 Downwind Marginal Distribution

The marginal distributions of particle position in the downwind direction have some features that can be exploited to achieve a good quality of fit across a wide variety of conditions with a single choice of function. The general shape of the distribution is consistent between cases: a sudden increase in particle landing density (fraction of particles per square meter), followed by a very long and gradual decline in particle landing density – a unimodal and right-skewed distribution. The tail of the distribution is very long, with nonzero density all the way to the end of the domain in the majority of cases. While the different particle species exhibit a wide range of average travel distance and tail length (as shown in Figure 6-10), the overall shape of the distribution remains generally consistent and shows a gradual change when comparing different particle species. An example of the general shape of the distribution is shown in Figure 6-12 for all particle species in a randomly selected case (64km/hr, High HRR, W:D 4, Forested Terrain):




Figure 6-12: KDE plot of particle landing distance separated by species. Lighter colour indicates higher particle ID number.

It is also worthwhile observing the shape of the distributions after taking the natural logarithm of the distance travelled downwind. We may find that fitting a distribution to the logarithm of distance is easier than fitting a distribution to the raw data. An example of the shape of the distribution of the log of particle travel distance is shown in Figure 6-13, showing a shape that resembles a normal distribution:




Figure 6-13: Example histogram and KDE plot of the natural log of particle travel distance for an example particle species.

Knowing the general shape of the distribution and log-transformed distribution, we can select a number of candidate distributions. For the raw data, the chosen candidate distributions are the Inverse Gaussian, Weibull, Rayleigh, and Pareto distributions. For the log-transformed data, the candidate distributions are the Normal and T-distributions. Each of the candidate distributions was fit to the data (either raw or log-transformed) separated into categories by particle species. Fitting was performed using a Maximum Likelihood Estimation (MLE) process using the Scipy.Stats library of functions in Python and the resulting distribution function parameters recorded.

Maximum Likelihood Estimation (MLE) is a well-established statistical method for finding the optimal parameters for a function fit to a given data set [92, 93]. Let us suppose we have a set of data $\mathbf{y} = y_1 \dots y_n$ and a probability density function that takes some parameters $\mathbf{w} = w_1 \dots w_k$. The likelihood function $L(\mathbf{y}|\mathbf{w})$ describes the likelihood that the data vector \mathbf{y} follows the probability density function with parameters in the vector \mathbf{w} . The likelihood function is defined as the product of the PDF for every data point in \mathbf{y} :

$$L(\mathbf{y}|\mathbf{w}) = \prod_{i=1}^n PDF(y_i, \mathbf{w}) \quad (6-1)$$

The process of maximum likelihood estimation attempts to find the values of parameters \mathbf{w} that maximize the likelihood value L for the given data. For some probability density functions an analytical solution for the optimal parameter vector $\hat{\mathbf{w}}$ can be found, such as in the case of a normal distribution. In other cases, numerical methods are required to estimate the optimal parameters. When numerical methods are used, it is possible that no optimal set of parameters

is found, or that a local maxima is found. It may be important to select a good initial guess for the parameters when using this technique.

Evaluating the goodness of fit for each distribution is a complicated process. There are a range of statistical measures of goodness of fit that we could use – such as the Kolmogorov-Smirnov distance, Cramér–von Mises criterion, or Anderson-Darling criterion. These statistical measures report a value that represents the difference between a theoretical cumulative distribution function (CDF) and an empirical cumulative distribution function (ECDF). In this context, the CDF is the cumulative density function of the fitted distribution and the ECDF is the cumulative density function of the simulated distribution of firebrands. While these measures are useful, they do not provide a complete picture of goodness of fit, and do not have a term to penalize overfitting. For our purposes, we will use three different methods to assess the goodness of fit of our distributions: the Akaike Information Criterion (AIC) [94], quantile-quantile plotting, and inspection of the plotted CDF and ECDF of our data – including an investigation of the Cramér–von Mises criterion [95].

The one-sample Cramér–von Mises criterion T is defined as the integral of the squared difference the CDF and ECDF weighted by the number of data points in the ECDF, as shown in Eqs. (6-2) and (6-3).

$$\omega^2 = \int_{-\infty}^{\infty} [F_N(x) - F(x)]^2 dF(x) \quad (6-2)$$

$$T = n\omega^2 \quad (6-3)$$

This provides a simple and intuitive measure of how well the CDF follows the empirical data. To illustrate this, two plots of Inverse Gaussian CDF functions with different Cramér–von Mises criteria are presented in Figure 6-14. The left plot, representing the fitted CDF and empirical data for particle species 41, shows a very good fit with a Cramér–von Mises criterion of 0.126. The right figure has a considerably larger Cramér–von Mises criterion value of 3.31, showing a much poorer quality of fit.




Figure 6-14: Example plots of the inverse gaussian CDF and ECDF of particle landing position for two particle types in the 64 km/hr, 1:1 W:D, high HRR, open field case. Particle 41 has a Cramér–von Mises criterion of 0.126, contrasting with Particle 24 with a criterion of 3.31.

The Akaike Information Criterion [94] can be used as a method of evaluating the relative performance of statistical models with a variable number of parameters. Naturally, a model with more parameters can be adjusted to better fit the data. However, a model with a large number of parameters can be “overfitted” to the data, leading to an overcomplicated model that performs poorly when used to make predictions. To avoid overfitting, the AIC takes into account the number of parameters fit by the model as well as the maximum likelihood estimator function:

$$AIC = 2k - 2\ln(\hat{L}) \quad (6-4)$$

Where \hat{L} is the maximum value of the likelihood function and k is the number of parameters. Given that we have used maximum likelihood estimation as the method for fitting our distributions, $\ln(\hat{L})$ can be calculated easily using our fitted probability density function as:

$$\ln(\hat{L}) = \sum_{i=1}^n \ln(PDF(y_i, \hat{\mathbf{w}})) \quad (6-5)$$

However, the AIC is not comparable between distributions that were fitted with the raw data and the log-transformed data without also transforming the PDF. This is possible for the normal distribution, producing the lognormal distribution, but is more difficult for the T-distribution. In general, we will only compare AIC values between PDFs that were fit with the same data (raw vs. log-transformed). Fitted distributions with smaller AIC values are favourable.

In addition to the AIC and Cramér–von Mises criterion, quantile-quantile plots (Q-Q plots) can be used to assess the goodness of fit of a distribution. These plots measure the relative positions of the quantiles of the fitted distribution against the empirical quantiles of the data. A perfect fit would have all points sit along a straight line at 45 degrees. The advantage of Q-Q plots over

numerical measures of goodness of fit like the AIC is that different areas of the distribution can be recognized on a Q-Q plot. For example, if a fitted distribution agrees well with the empirical data everywhere except at one extreme, the Q-Q plot will show exactly where the fitted distribution starts to disagree with the empirical data.

Finally, inspection of the fitted CDF against the ECDF obtained from the data allows for a straightforward visual understanding of the quality of fit. Ideally, the fitted CDF and ECDF should overlap completely, but in cases where they do not, visual inspection can show both how much discrepancy there is between the fitted CDF and the ECDF and where the discrepancy occurs. Combined with the Cramer Von Mises criterion (which represents the total area between the CDF and ECDF), this should provide a final means of deciding which of the candidate distributions is the best fit for the data.

A table summarizing the quality of fit for each candidate distribution across all scenarios is presented in Table 6-4.

Candidate Function	Mean AIC	Mean Cramér-von Mises criterion	Median Cramér-von Mises criterion	Failure To Fit Pct.
Log T	11087.056*	1.8241	0.6034	0%
Log Normal	11584.934*	5.9874	2.3077	0%
Inverse Gaussian	32531.673	3.8677	0.4993	0%
Weibull	35487.212	16.791	1.5540	9.94%
Rayleigh	36302.100	75.916	26.672	0%
Pareto	32613.122	6.6755	2.0026	0%

Table 6-4: Akaike Information Criterion values for the selected candidate functions. *AIC values of log transformed data

Based on the results presented in Table 1, we can evaluate the relative performance of each candidate function. For functions fit to log-transformed data, the log T distribution performs better than the log normal distribution, with a superior AIC and Cramér-von Mises criterion value. For functions fit to non-transformed data, the inverse gaussian function has the best AIC and Cramér-von Mises criterion by far. The Weibull, Rayleigh, and Pareto distributions do not have good enough performance to be worth considering for further evaluation. This leaves two possible candidate functions: log T and Inverse Gaussian. To decide between the log T and Inverse Gaussian functions, we can consider the shape of the Q-Q plots for these functions across a variety particle species and visually inspect the differences between the CDF of each function. An example Q-Q plot for these functions is presented in Figure 6-15:





Figure 6-15: Quantile-Quantile plots for the Inverse Gaussian and log T distributions. 64km/hr 4:1 W:D High HRR Forested case.

Based on an inspection of the Q-Q plots and visual inspection of the CDF against empirical data, both functions are a very close fit to the simulated distribution. However, we can observe

that there are some consistent trends for each function that lead to an under- or over-prediction of particle density in certain areas. The inverse gaussian function tends to perform very well in the region closer to the fire before gradually under-predicting the quantity of firebrands in more distant regions, performing worse at the far end of the domain. The log T distribution over-predicts the quantity of firebrands at each end of the distribution, and slightly under-predicts the quantity of firebrands landing in the middle. As each distribution is so close in performance based on the previously discussed criteria, it is difficult to choose between them – but we may prefer that our errors are concentrated in one region or another. Due to the very close margin of performance between these two distributions, both will be compared and evaluated for use in the final model. This description of particle landing density in the downwind direction using the log-T and Inverse Gaussian distributions may be a novel finding – these functions appear to fit the simulated particle landing data better than some existing probability density functions used in literature (e.g. log-normal).

6.5.2 Crosswind Marginal Distribution

The marginal distribution in the crosswind axis is considerably more difficult to work with than that of the downwind axis. While the downwind distribution closely resembled a number of existing probability density functions and tended to have the same shape regardless of the conditions of the simulation, this is not true for the crosswind distribution. Different wind speeds and W:D ratios produce very different marginal distribution shapes in the crosswind direction, meaning that we may not be able to find a convenient pre-existing probability density function that accurately describes our data. This means the process of fitting a function to the crosswind marginal distribution may require different techniques to those that were used for the downwind distribution. To begin, we should return to the kernel density estimate plots described in the data visualization section of this chapter. A selection of KDE plots is presented in Figure 6-16, illustrating the change in the shape of the distribution as W:D ratio increases:






Figure 6-16: KDE plots illustrating the change in distribution shape across different $W:D$ ratios for a selection of open field cases.

A clear trend can be seen in the shape of the distribution as $W:D$ increases for open field cases, with higher $W:D$ ratios producing a bimodal shape in the kernel density estimate plot. However, for cases in forested areas, the shape can become less well-defined, with more noise and smaller peaks:

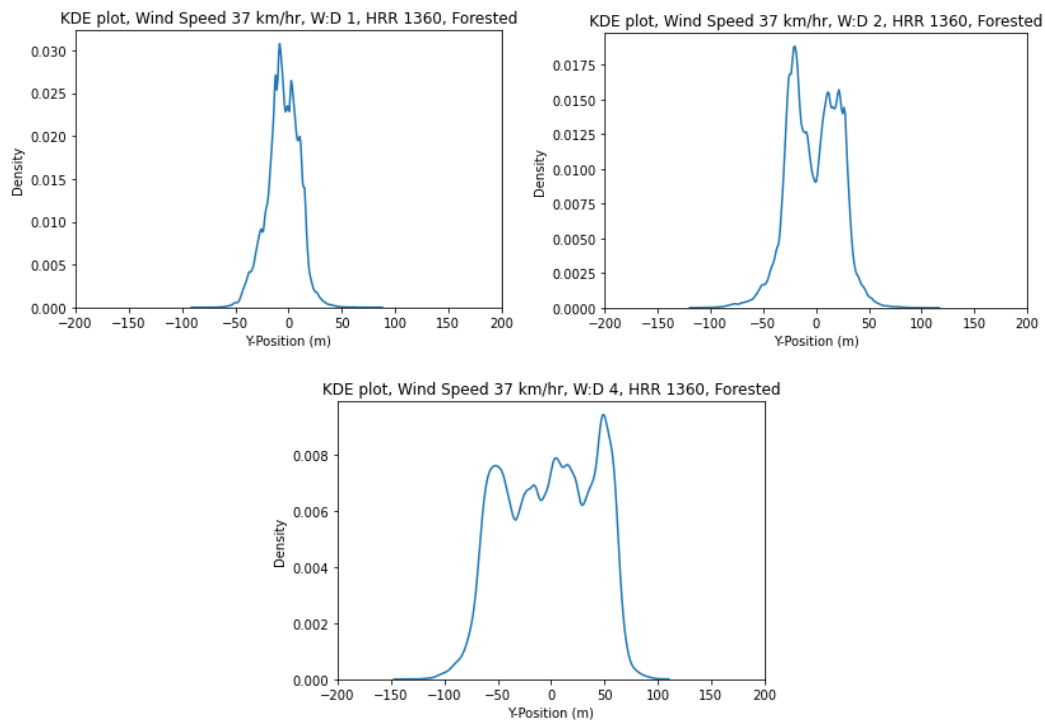



Figure 6-17: KDE plots illustrating the change in distribution shape across different $W:D$ ratios for a selection of forested cases.

To describe the variety of shapes in these marginal distributions, we will need to find or develop a function that can transition between a single-peaked (unimodal) and multi-peaked (bimodal) shape. We can note some useful properties of the crosswind marginal distributions: they tend to be roughly symmetrical, the edges of the distribution show steep increases and decreases in

probability density, and the width of the distribution is usually not much greater than the width of the fire that generated the distribution.

Keeping these properties in mind, we will use an iterative approach to develop a function that can describe the variety of distributions obtained from our simulations. The first potential function that may approximate the distributions is a Gaussian Mixture: a weighted sum of two gaussian distributions with different values of μ and σ . This function can smoothly transition between single-peaked and double-peaked by changing the values of μ for each component (close values produce a single peak, different values produce a double-peak) and should be close to symmetrical for our data. Testing this function, we obtain a good approximation of the crosswind marginal for some cases, but the quality of the fit is not consistent:

prt5_params_100cm_37kmhr_1w_lowheat Gaussian Mixture fit - Scaled



prt5_params_100cm_54kmhr_4w_lowheat Gaussian Mixture fit - Scaled




Figure 6-18: Plots of a Gaussian Mixture fit to the crosswind marginal distribution showing varying quality of fit. top plot shows a good quality of fit, the bottom plot shows poor quality of fit.

We can quantify the quality of fit for this function using the Cramér-von Mises criterion, as was used in the process of fitting functions to the downwind marginal distribution. Using this metric for quality of fit, we can see that the gaussian mixture function performs reasonably well for certain cases but does not perform well in general. A summary of the performance of this model is presented in Table 2:

	Open Field		Forested	
W:D	Mean CVM	Median CVM	Mean CVM	Median CVM
1:1	113.34	11.279	429.43	219.07
2:1	15.501	14.311	79.663	40.337
4:1	46.005	36.081	131.69	41.680

Table 6-5: Performance summary of the Gaussian Mixture model.

Based on the data presented in Table 2, and by inspection of plots of the Gaussian Mixture probability density functions overlaid with histograms of the raw data, we can see where the Gaussian Mixture functions tend to perform well and where they fail. The Gaussian Mixture functions perform best in cases that are single-peaked, or where the distance between peaks is relatively small and the transition is gradual. As the distance between peaks increases, the function becomes less and less accurate. The Gaussian Mixture function also performs worse in forested cases than in open field cases. This is likely due to the “noisy” regions present in the centre of the crosswind marginal distributions of many of the forested cases (such as in the 4:1 WD case shown in Figure 6-17).

Iterating on the Gaussian Mixture function, we can introduce an additional term to each component of the Gaussian Mixture to make a new function – a Skewnormal mixture. The Skewnormal distribution is a modified normal distribution with a probability density function defined as:

$$PDF_{SN}(x, \alpha) = 2\varphi(x)\phi(\alpha x) \quad (6-6)$$

And the Skewnormal mixture probability density function is defined as:

$$Mixed\ PDF = A_1 PDF_{SN}\left(\frac{x + \varepsilon_1}{\omega_1}, \alpha_1\right) + A_2 PDF_{SN}\left(\frac{x + \varepsilon_2}{\omega_2}, \alpha_2\right) \quad (6-7)$$

Where $\varphi(x)$ is the probability density function of the normal distribution, and $\Phi(\alpha x)$ is the cumulative density function of the normal distribution. The value of α determines the skewness of the distribution – a positive value of α shifts the peak of the probability density function to the right and extends the tail of the distribution rightward. The terms A , ε , ω determine the relative magnitude, position, and scale of each component of the mixed function. Using the Skewnormal distribution, we can achieve a better quality of fit for wider distributions that have steep edges – a common feature of observed distributions in 4:1 WD scenarios. A comparison of the quality of fit between a Gaussian Mixture and a Skewnormal Mixture function is presented in Figure 6-19:




Figure 6-19: Comparison of quality of fit between Gaussian Mixture and Skewnormal Mixture functions.

Using the Cramér-von Mises criterion to evaluate the quality of fit for the Skewnormal Mixture function, we can see a dramatic improvement in quality of fit in every category:

	Open Field		Forested	
W:D	Mean CVM	Median CVM	Mean CVM	Median CVM
1:1	20.853	9.3864	26.140	13.563
2:1	2.1545	1.2127	17.511	8.6730
4:1	5.2319	2.7601	19.855	14.611

Table 3: Performance summary of the Skewnormal Mixture model.

While this function performs very well in the majority of cases, there is an underlying problem that has to be addressed before the function could be used in practice. Each component of the

Skewnormal mixture has slightly different parameters, producing an asymmetrical result – but the initial conditions of each scenario do not have any asymmetrical properties. The function that we choose to use as an approximation of the crosswind marginal distribution should be symmetrical to reflect the symmetry of the initial conditions. There are a few methods that we could use to produce a symmetrical probability density function from the asymmetrical function we have fit, possibly in combination with each other. One option is to average the function with a reflection of itself around the centre line of the domain. Another option is to define a new function where the parameters of each component are related to each other – linking the values of A , α , ε , ω in Eq. (6-7). A third option is to alter the data set so that it is perfectly symmetrical (summing the data with its own reflection around the centre line) and fitting the Skewnormal mixture function to the symmetrical data set.

Using a modified Skewnormal mixture function with linked values of A , α , ε , ω and fitting to a symmetrical data set, we obtain a high quality of fit with a very similar shape to the original Skewnormal mixture function, shown in Figure 6-20:

$$Mixed\ PDF = \frac{1}{2}PDF_{SN}\left(\frac{x + \varepsilon}{\omega}, \alpha\right) + \frac{1}{2}PDF_{SN}\left(\frac{x - \varepsilon}{\omega}, -\alpha\right) \quad 6-8$$




Figure 6-20: Plot of a modified Skewnormal mixture function fit to a symmetrical data set.

The modified function fit to symmetrical data performs very well in most cases, with only a few outlier cases where the quality of fit was poor. In general, it has comparable performance to the original Skewnormal mixture function while having a symmetrical output.

	Open Field		Forested	
W:D	Mean CVM	Median CVM	Mean CVM	Median CVM
1:1	8.9698	3.0883	17.480	2.5953
2:1	3.6130	1.0264	90.914	25.465
4:1	0.7795	0.6567	30.135	10.846

Table 4: Performance summary of the modified Skewnormal Mixture model.

As this modified Skewnormal function has comparable quality of fit to the original function but has the symmetrical properties we want for our final model, this function was selected as the best option to fit for the crosswind marginal distribution. This function for the distribution of particle landing density in the crosswind axis is certainly a novel finding.

6.6 Describing the Joint Distribution

Now that a function has been selected for each marginal distribution, we must find a function that can describe the interaction between these functions (the “dependence structure”) and produce a realistic two-dimensional joint probability density function.

6.6.1 Copula Method

There are multiple approaches that could be used to try and find such a function – we will consider two approaches. The first approach is to find the *copula* of the distribution. This technique relies on Sklar’s Theorem [96]: once the marginal distributions are known, the underlying dependence structure can be obtained by transforming the data so that the marginal distributions are uniform. The distribution that remains after the marginal distributions have been transformed to be uniform is the dependence structure. This technique is used in a variety of fields, but there is no guarantee that an empirically obtained copula will be described by an easily expressed function. In our case, we can visualize the copula of our data and compare it with a kernel density estimate plot to gain some insight into the shape of the dependence structure, as in Figure 6-21. In this figure, the marginal distributions of the KDE plots on the left are transformed to a uniform distribution (i.e. a flat and constant probability density) and normalized onto the range $[0,1]$ to produce the copula plots on the right.

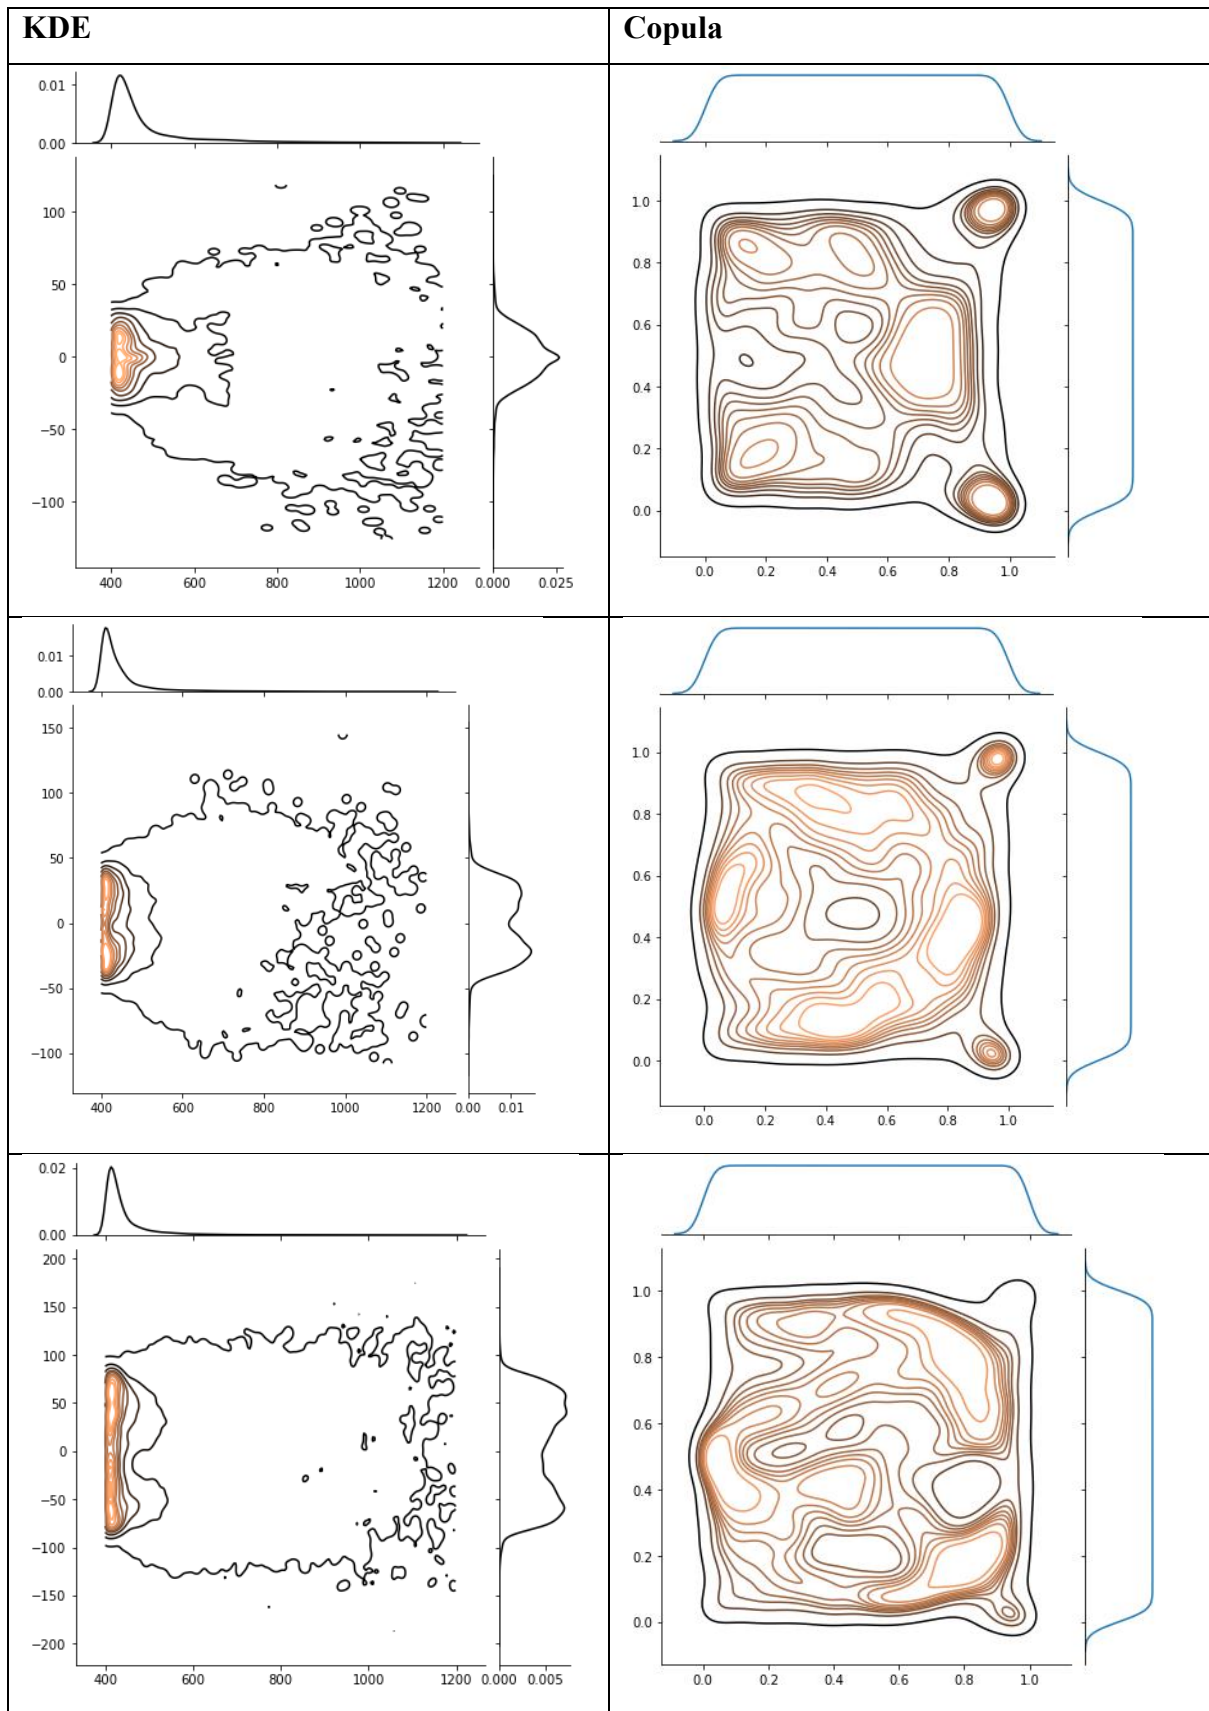


Figure 6-21: Comparison of a Kernel Density Estimate plot and the underlying Copula. **Top:** 64 km/hr, $W:D = 1$, High HRR, Open Field. **Middle:** 46km/hr, $W:D = 2$, High HRR, Open Field. **Bottom:** 46km/hr, $W:D = 4$, High HRR, Open Field.

Although the copulas shown in Figure 6-21 are quite complex, these visualizations reveal a trend in our data that could not be clearly seen otherwise. For the set of copulas generated from open field scenarios, there is a clear trend in the location of the region of lowest probability density in the copula. The location of the low-density region depends strongly on the W:D ratio of the scenario. Scenarios with a W:D of 1 produce a copula with a large low-density region close to the fire. As the W:D ratio increases, the low-density region moves rightward – with a W:D of 2 placing the low-density region in the centre of the domain, and a W:D ratio of 4 placing the low-density region close to the right edge of the domain.

HRR and wind speed do not seem to have a noticeable effect on the overall shape of the copula in open field cases. This indicates that the effect of these parameters is primarily reflected in the shape of the marginal distributions, and that it does not alter the shape of the dependence structure of the distribution.

The copulas obtained from forested scenarios tend to be much more complex than those obtained from open field cases. In general, these copulas share some characteristics with each other, but do not show the same trends as the open field cases. These copulas tend to exhibit a large low-density area at the right edge of the domain bordered by two higher density areas at the corners. Cases with a smaller W:D show this behaviour more strongly than cases with larger W:D ratios, but it is present in almost all forested cases.

Due to the high level of complexity of the copulas obtained from both open and forested scenarios, it is unlikely that there is a convenient function that can adequately approximate the observed behaviour. However, qualitative observations about the copulas may inform our decisions in describing the joint probability density function by other methods.

6.6.2 Conditional Probability Method

The second method for describing the joint distribution of the data is based on the analysis of conditional probability density functions. By dividing the distribution into segments of known width along the downwind axis, we can fit a function to each segment and find a conditional probability density function for each segment along the downwind axis. The joint distribution can then be approximated by this collection of conditional probability density functions.

To describe this method more formally, let $f_{XY}(x, y)$ be the two-dimensional joint probability density function of our distribution, and let the function $f_X(x)$ be the marginal probability density function in the x-axis (the downwind axis). The function $f_{Y|X}(y|x)$ is the conditional

probability density function of our distribution at a particular point x . Using a theorem described in the textbook by Park [97], it can be shown that:

$$f_{Y|X}(y|x) = \frac{f_{XY}(x,y)}{f_X(x)} \quad (6-9)$$

Given that we have a good approximation for the marginal distribution $f_X(x)$, if we can find a good approximation for $f_{Y|X}(y|x)$ for all x then we can describe the joint distribution based on these functions. We can approximate $f_{Y|X}(y|x)$ for a particular x by finding a function that describes the distribution in an interval $[a, b]$ where $a < x < b$. To find $f_{Y|X}(y|x)$ on $[a, b]$ we can use the same process we used to describe the crosswind marginal distribution but applied to each segment – plot a histogram of the points, and fit a curve to the data.

In practice, there is a restriction on how narrow the interval $[a, b]$ can be. Our data is a collection of points, not a continuous surface. This means we need to choose an interval width large enough to capture a reasonable number of points for our curve-fitting process to work – fewer points will produce a worse quality of fit. This means that as we approach the far downwind areas of the domain where very few particles land, the interval width will grow to be very large. This makes it hard to capture fine details about the shape of the distribution in this region.

A division of the domain into nine segments of variable width was chosen. Between 400 and 500m, segments are 20m wide. From 500 to 600m, segments are 50m wide. One segment spans the distance from 600 to 800m, and another from 800 to 1200m. This is represented in Figure 6-22, overlaid with a kernel density estimate plot:




Figure 6-22: Division of the domain into segments.

Using these segments and the methodology previously used for fitting the crosswind marginal distribution, we can obtain a set of nine histograms and observe how well the mirrored Skewnormal function fits. An example is shown in Figure 6-23:

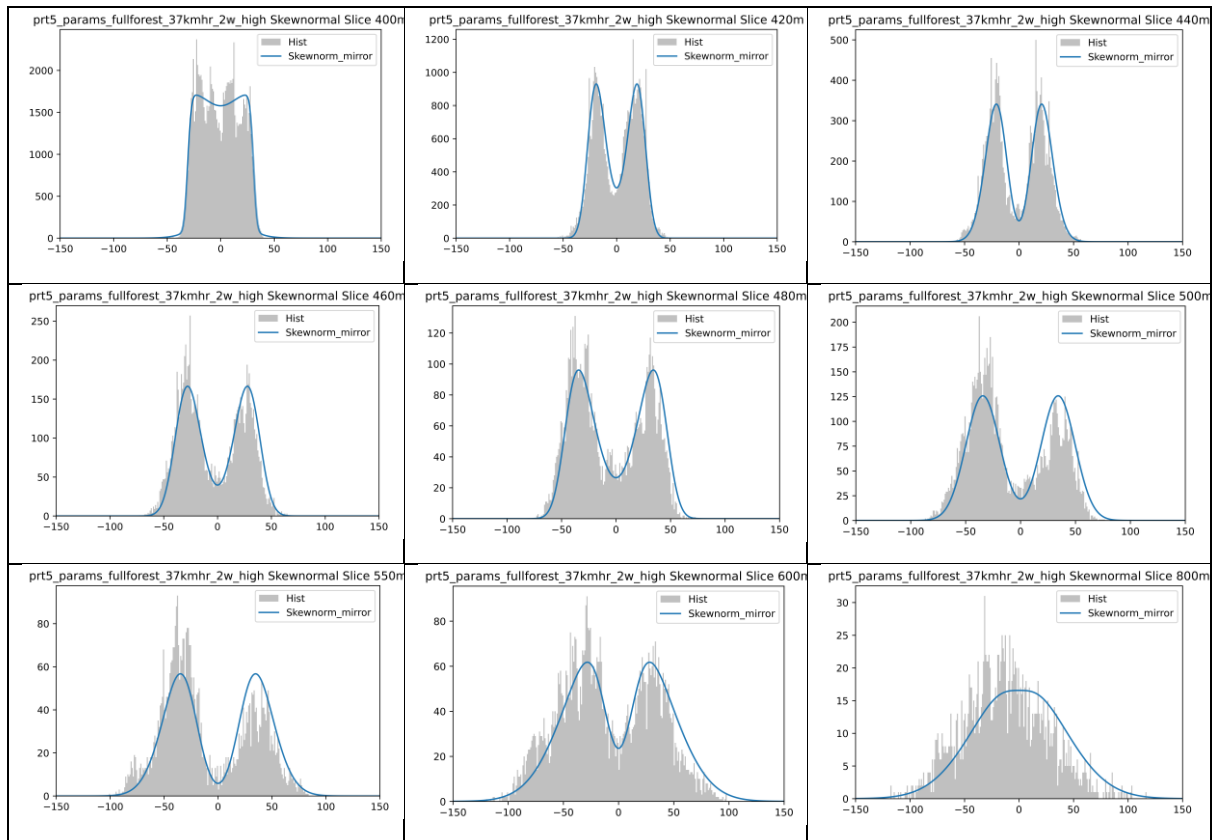


Figure 6-23: Plots of conditional distributions for each segment. Title indicates the distance at which the segment begins.

As shown in Figure 6-23, the mirrored Skewnormal function fits each of the conditional probability distributions relatively well for this scenario. The quality of fit is highest for segments in the region beyond 420m and before 800m. The segment closest to the fire (400-420m) shows a large, flat distribution that does not have clear, smooth peaks like most of the conditional distributions in regions further downwind. The final segment (800-1200m) also lacks well-defined peaks, somewhat resembling a normal distribution.

Based on these plots, and similar plots generated from all other scenarios, we can conclude that the mirrored Skewnormal function is adequate to describe the conditional probability distributions throughout the domain. This means that the joint distribution $f_{XY}(x, y)$ can be described as:

$$f_{XY}(x, y) = f_X(x)f_{Y|X}(y|x) \quad (6-10)$$

Where:

$$f_{Y|X}(y|x) = \frac{1}{2}PDF_{SN}\left(\frac{x + \varepsilon}{\omega}, \alpha\right) + \frac{1}{2}PDF_{SN}\left(\frac{x - \varepsilon}{\omega}, -\alpha\right) \quad (6-11)$$

The values of α , ε , and ω (skew parameter, location, and scale) are functions of x and the initial conditions of the scenario. The next step is to characterize how these parameters vary with x , and then finally to relate the values of these parameters at $x = 400$ to the initial conditions of the scenario.

It is important to note that while the curve fitting process produced a good quality of fit for the majority of cases and segment locations, some segments produced a low quality of fit and had their data points excluded from the following fitting processes. Of 423 segments, six were excluded as outliers for producing implausibly high values of skewness relative to similarly located segments.

Plotting the skew parameter term α over downwind distance x reveals a clear trend. The value of α depends strongly on x , following a hyperbolic curve of the form:

$$\alpha = \frac{A}{x - K} + H \quad (6-12)$$

This can be observed in both forested and non-forested cases, and for all variations of W:D, speed, and HRR. This is shown in Figure 6-24:




Figure 6-24: Plot of the Skewness term over distance for forested and non-forested scenarios. Hyperbolic curves fit to the data show some impact of forest on skewness.

It can be seen from Figure 6-24 that the effect of HRR on skewness is minimal in open field scenarios. In forested scenarios, there is more variation between cases that could be attributed to differences in HRR, but it is still considerably less than the effect of other parameters. The effect of wind speed on skewness is also relatively small, with no consistent increase in skewness observed between the 37km/hr cases and 64km/hr cases. The most significant simulation parameter affecting skewness is W:D ratio, affecting the magnitude of the hyperbolic curve dramatically. This is shown in Figure 6-25:

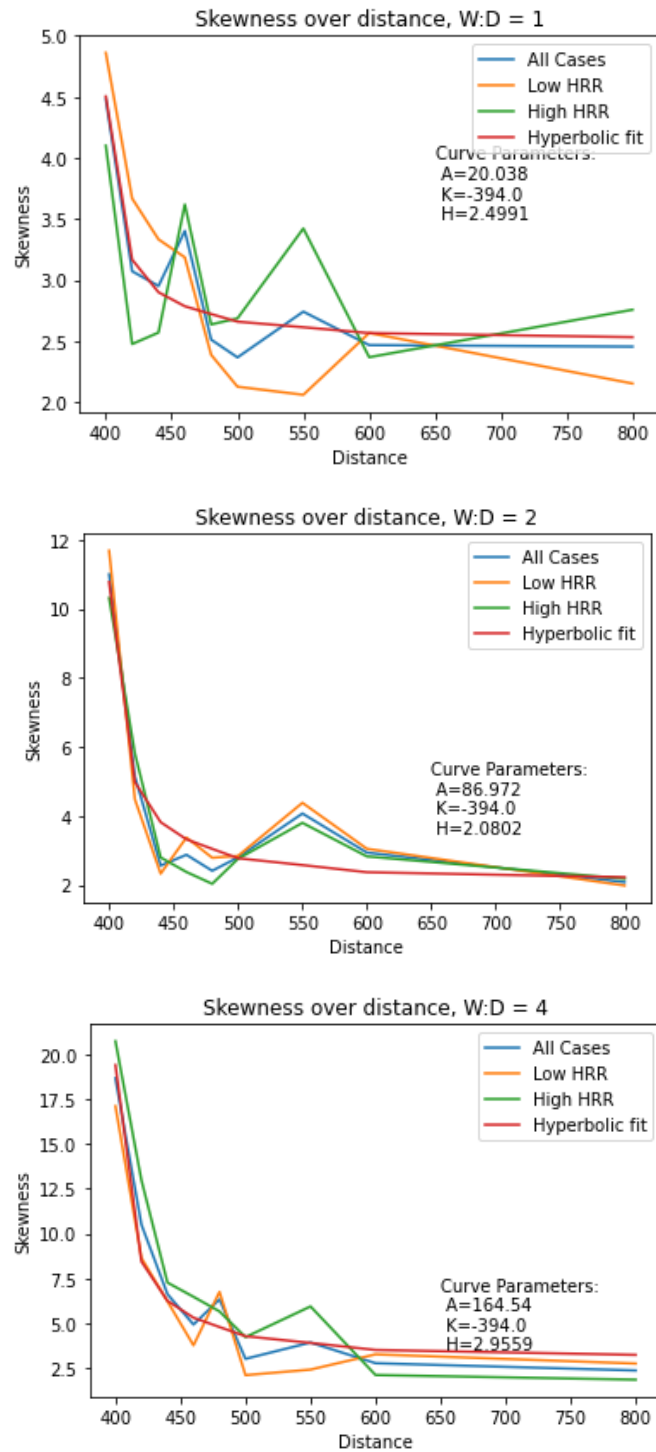


Figure 6-25: Hyperbolic curves fit to skewness for cases separated by $W:D$ ratio. Higher $W:D$ cases have considerably higher skewness in the regions near the fire.

The skewness term of the distribution of firebrand particles at $x = 400$ is almost four times higher in cases with a $W:D$ ratio of 4 than in cases with a $W:D$ ratio of 1. Skewness in the regions far downwind of the fire are almost identical between cases regardless of $W:D$ ratio,

asymptotically approaching a value between 2 and 3. Combined with previous observations about the effect of forest conditions, HRR and wind speed on skewness, it is likely that skewness could be reasonably approximated as a function of three parameters: distance, W:D ratio, and presence or absence of forest.

The location term ε in equation (6-11) is primarily influenced by the real width of the fire – a function of both wind speed and the W:D ratio of the fire. However, the change in the location term over distance does not seem to follow any consistent trend with any other parameter. The location term seems to increase and decrease at random between each segment. Due to the seemingly random fluctuations of this value, this term may be best approximated as a constant in x determined by the initial conditions of the simulation.




Figure 6-26: Dependence of the Location term at $X=400$ on real flame width.

Finally, the scale term ω can be seen to vary significantly depending on certain parameters and downwind distance. As with skewness, the W:D ratio and presence or absence of forest in the area downwind of the fire are important parameters for this term. HRR seems to have relatively little impact on this term in open field scenarios but has a considerable impact in forested scenarios.

Interestingly, while the scale parameter is clearly impacted by distance, W:D ratio, and forest conditions, it does not follow a single simple curve like skewness. The shape of the curve of the scale parameter over distance changes depending on the W:D ratio, from almost linear at

low W:D to a more complex shape at higher W:D ratios. This change in shape is illustrated in Figure 6-27:

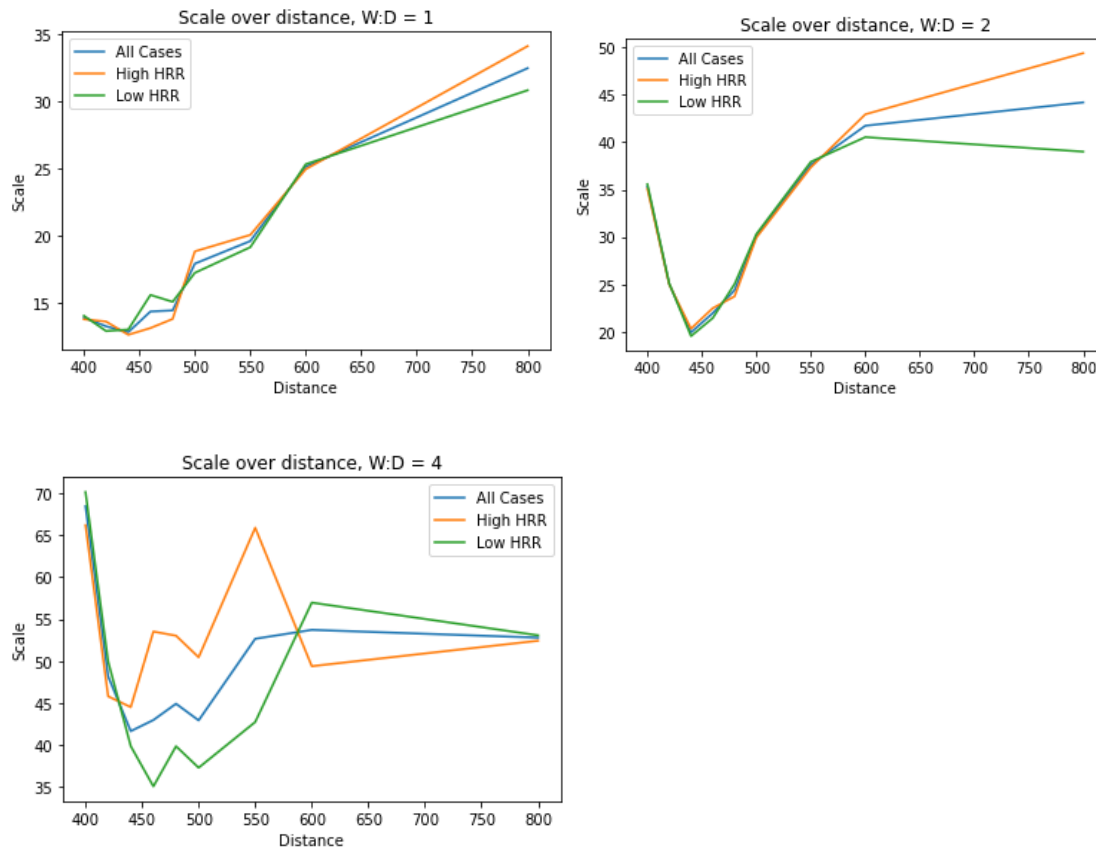


Figure 6-27: Change in the shape of Scale over Distance curves for different W:D ratios.

As the shape of the curve changes between W:D ratios, it is likely to be difficult to find a simple function that captures this behaviour accurately. For the purposes of our model, we may fit a function for each W:D case separately and interpolate between them to find an approximation of the scale term for intermediate W:D values.

6.7 Model Fitting

We now have a complete picture of all the equations necessary to describe the joint distributions we have obtained from simulation. Equations (6-9) and (6-10) describe how the downwind marginal probability density functions $f_X(x)$ and conditional crosswind probability density functions $f_{Y|X}(y|x)$ can be combined to produce the joint distribution. The downwind marginal probability density function $f_X(x)$ has been found to be closely approximated by a log T or Inverse Gaussian function whose shape depends on particle geometry, wind speed, W:D ratio, forest conditions and HRR. The crosswind conditional

probability distribution $f_{Y|X}(y|x)$ is a more complicated function, with many terms that need to be approximated or interpolated.

The final step in this modelling process is to relate each of the terms in these equations to the underlying physical conditions of our simulations. For the downwind marginal probability density function, this might be achieved by implementing a linear model for each particle type and treating the overall downwind marginal PDF as the sum of the PDFs of each particle type.

Before generating such a model, we should investigate how the parameters of the log T and Inverse Gaussian distributions relate to the physical conditions of the simulation graphically. For preliminary model development, we will focus on the Inverse Gaussian function, as it is mathematically simpler. Unlike the crosswind conditional probability distributions, the parameters of this function depend strongly on the mass and geometry of each particle. This makes it difficult to generalize about trends in each parameter, as the trends for one type of particle may not hold for another. For example, the μ parameter for the Inverse Gaussian distribution varies significantly according to particle type. When we split the data according to particle geometry, we can clearly see the effect of HRR, forest conditions, and W:D ratio on μ depend on both the geometry of the particle and its mass:

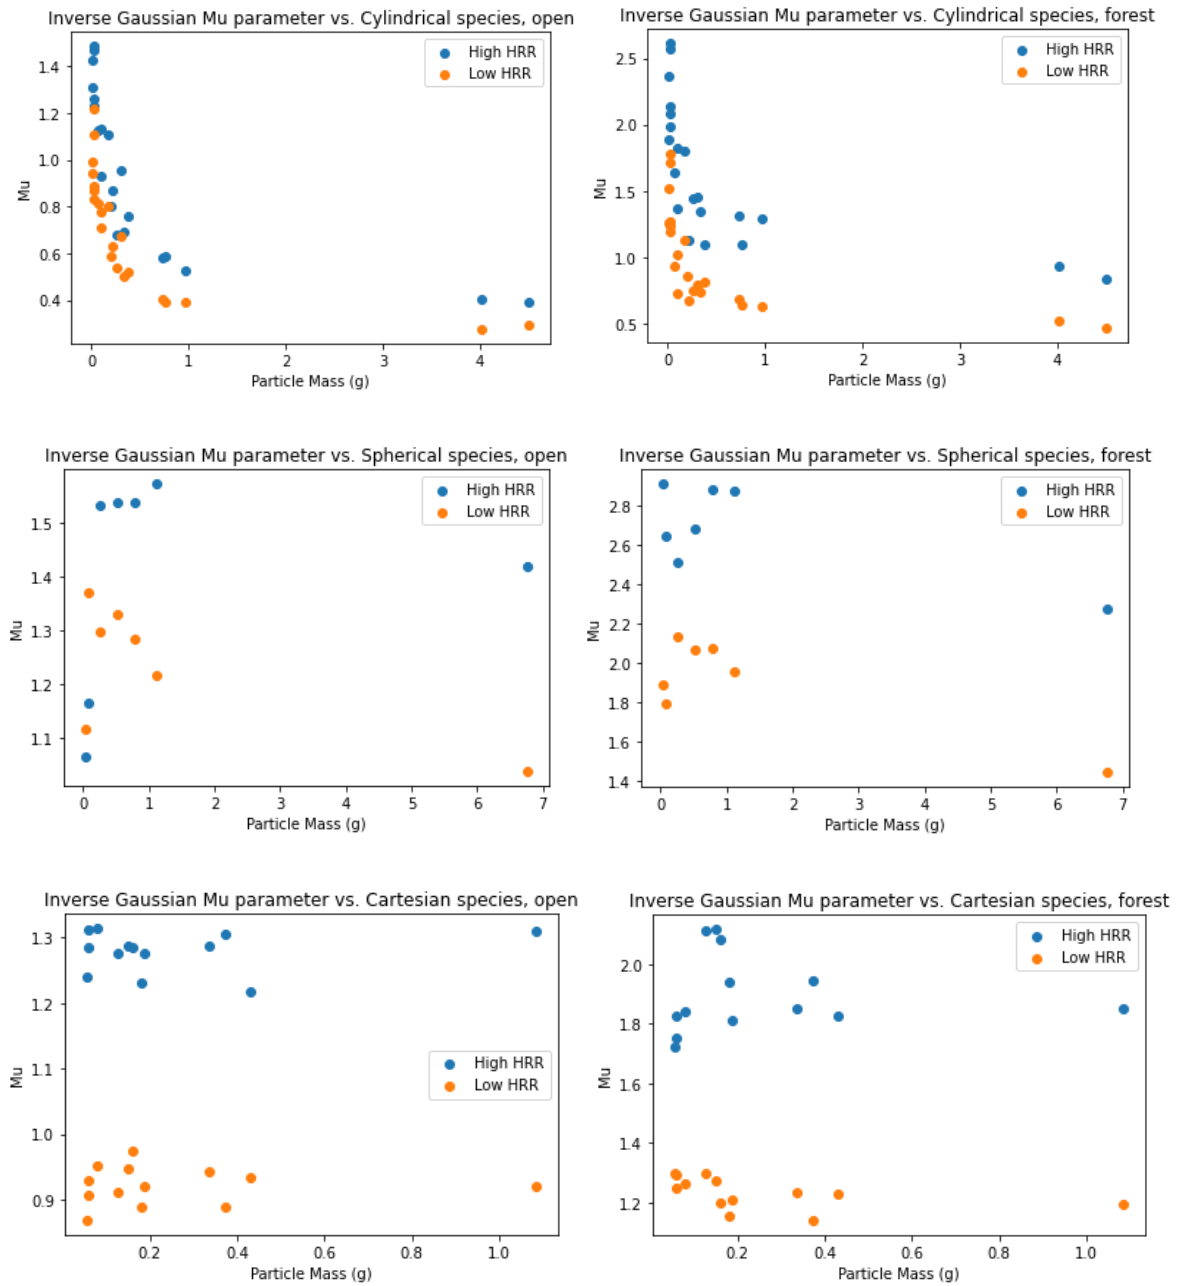


Figure 6-28: Effect of HRR on μ parameter of the Inverse Gaussian function for different particle masses and geometries, separated by forest condition.

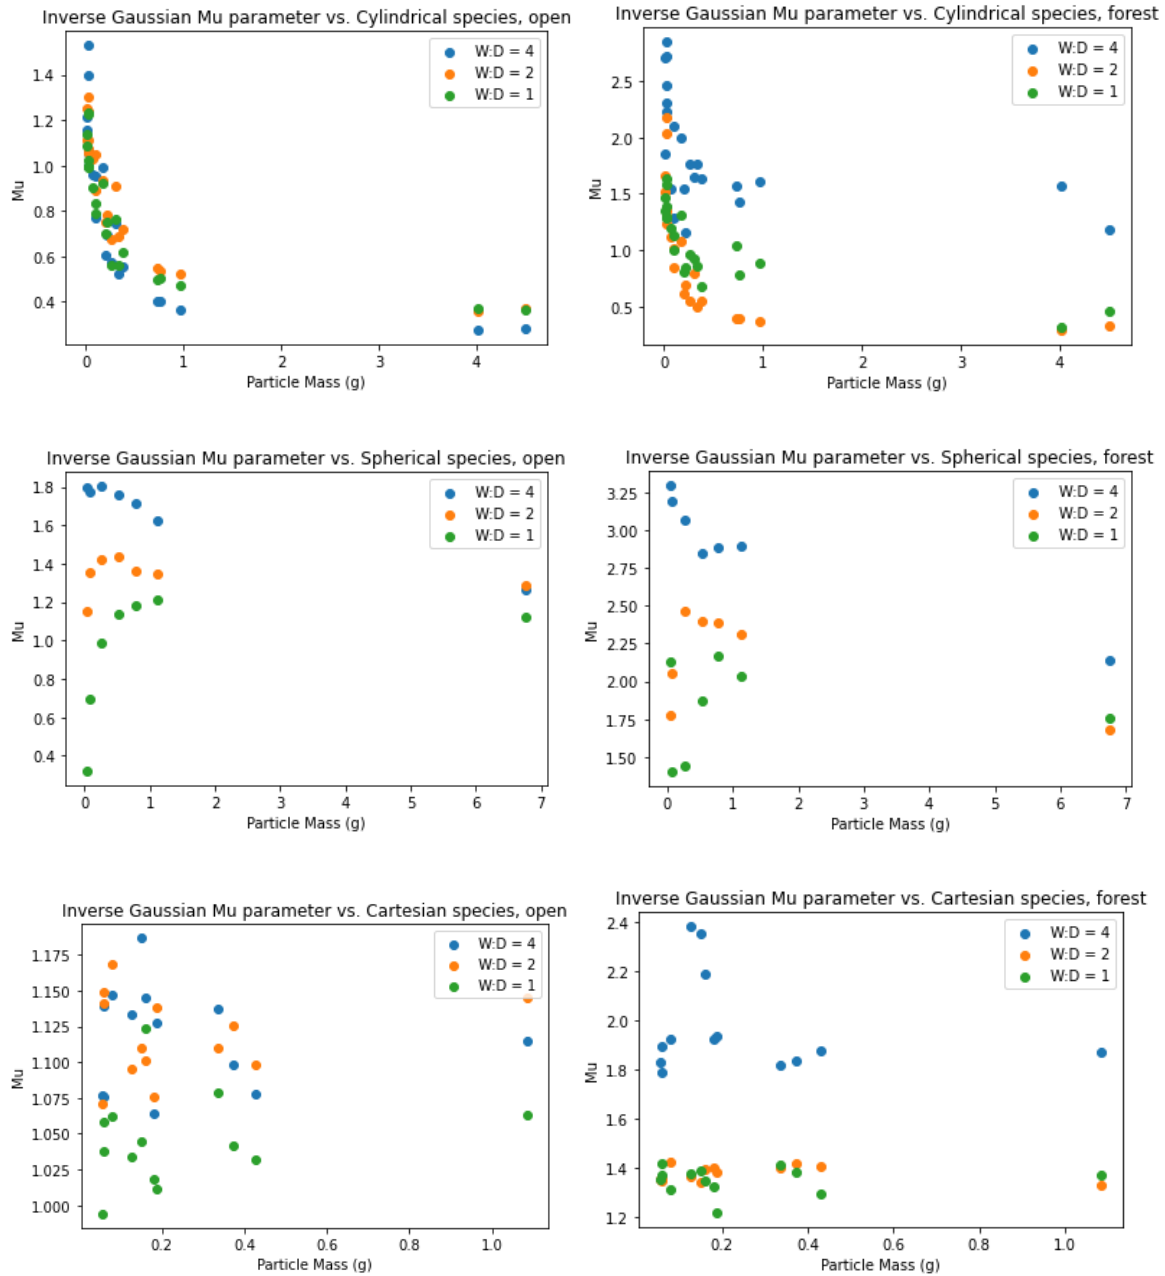


Figure 6-29: Effect of W:D ratio on μ parameter of the Inverse Gaussian function for different particle types and forest conditions.

Due to the complexity of these relationships, preliminary model fitting may need to rely on some simplified approximations of the real behaviour of each parameter. However, certain relationships do seem to be apparent from initial observation: In almost all cases, higher HRR produces a higher value of μ . Forested scenarios have a greater range of μ values but tend to follow the same trends as non-forested scenarios. Based on these results, we can fit a curve to the parameters of the Inverse Gaussian function based on particle mass and the physical conditions of the scenario.

The μ and scale parameters of the Inverse Gaussian function seem follow a linear relationship with the log of particle mass for most particle types. By fitting a linear function to this curve for each particle geometry under each set of initial conditions, we can generate a table of parameters that can approximate the shape of the Inverse Gaussian function for each particle species in each scenario.

Inverse Gaussian Mu parameter vs. ln(Mass), Cylindrical species, HRR low, open




Figure 6-30: Linear relationship between μ and $\ln(\text{mass})$ for cylindrical particles.

This table of linear relationships is sufficient to complete a preliminary, novel model of particle landing density. We can now state a formulation of our new model as a series of equations and a table of parameters:

$$f_{XY}(x, y) = f_X(x)f_{Y|X}(y|x) \quad (6-13)$$

Where:

$$f_{Y|X}(y|x) = \frac{1}{2}PDF_{SN}\left(\frac{y + \varepsilon}{\omega}, \alpha(x)\right) + \frac{1}{2}PDF_{SN}\left(\frac{y - \varepsilon}{\omega}, -\alpha(x)\right) \quad (6-14)$$

$$\alpha(x) = \frac{A}{x} + H \quad (6-15)$$

Parameters ε and ω are approximated using a combination of linear functions of x depending on the initial conditions of the simulation. Parameters A and H depend only on the initial conditions of the simulation and are not functions of x .

The equation for $f_X(x)$ is expressed as the sum of the PDFs for each particle species p weighted by their relative frequency W_p :

$$f_X(x) = \sum_{p=0}^n \frac{w_p}{n} PDF_{Inv.G}(x, \mu_p, \omega_p) \quad (6-16)$$

$$PDF_{Inv.G}(x, \mu_p, \omega_p) = \frac{1}{\sqrt{2\pi(\frac{x}{\omega_p})^3}} \exp\left(-\frac{\left(\frac{x}{\omega_p} - \mu_p\right)^2}{2\frac{x}{\omega_p}\mu_p}\right) \quad (6-17)$$

Where the parameters μ_p and ω_p are functions of the mass and geometry of the particle species p and the conditions of the scenario. These functions are generally of the form:

$$\mu_p = B * \ln(m_p) + C \quad (6-18)$$

$$\omega_p = D * \ln(m_p) + E \quad (6-19)$$

Where the values of B, C, D, and E are selected from a table of values for each type of particle geometry under a certain set of initial conditions, and m_p is the mass of a given particle.

We can now compare the distribution of particles obtained from simulation with the distribution of particles predicted by this preliminary model:




Figure 6-31: Comparison of the simulated distribution of particles (left) and the statistical model of particle landing (right) for the 46km/hr, 4:1 WD, High HRR open case.

Visual comparison of these distributions shows that the statistical model performs quite well, producing an output with a similar plot of probability density contours to the simulation results. However, some areas for improvement can be identified. The statistical model slightly under-

predicts the width of the distribution in the crosswind axis at long distances, and slightly over-predicts width at close distances. This is a sign that the approximation of the parameters in $f_{Y|X}(y|x)$ could be improved in this case. It may be useful to split the functions that describe the parameters of $f_{Y|X}(y|x)$ into separate cases based on particle geometry – a separate function for cylindrical, spherical, and cartesian firebrands. However, splitting this function into the sum of many components may lead to difficulties in producing a good quality of fit in the far downwind region – as few particles land in this area, further dividing this population by species would lead to a very sparse dataset to fit a function to. Fitting the function $f_{Y|X}(y|x)$ to these sparse per-particle datasets in the far downwind region would likely produce a low quality of fit, and not improve model performance overall.

Development of a quantitative measure of model performance is necessary to evaluate the performance of the current model and future modifications to the model. This is investigated further in the next chapter – Model Description and Performance Evaluation.

Iterative development of this model may yield further improvements in speed and accuracy. Future research may produce better functions for approximating the parameters of $f_X(x)$ and $f_{Y|X}(y|x)$ for different particle species. Inspection of the quality of fit for each parameter may reveal some potential areas for further study to improve the performance of the model. Future development may also test the performance of the Log-T distribution as an alternative to the Inverse Gaussian distribution for the PDF of $f_X(x)$.

Chapter 7 Model Description and Performance Evaluation

Contents

Introduction.....	185
Model Description	185
Downwind Marginal Distribution.....	185
Crosswind Conditional Probability Distribution	18680
Model Performance.....	188
Performance Measurement Statistics.....	188
Overall Performance Statistics and Summary	194
Model Performance in Specific Scenarios	198
Open Field Cases	199
Forested Cases	205
Conclusion	205

7.1 Introduction

Having developed a simplified model of firebrand distribution based on the statistical analysis described in the previous chapter, it is necessary to present the model as a practical set of functions that can be used to produce a distribution of particles based on a set of input parameters, and to evaluate the performance of this model relative to the distribution obtained from CFD simulation. This chapter will present the model as a set of mathematical functions and compare the output of these functions to data obtained from CFD simulation. The quality of the model output will be evaluated using both qualitative and quantitative techniques.

7.2 Model Description

The four input parameters required to produce a firebrand landing distribution using this model are:

1. Domain type (Forest/Open Field)
2. Wind Speed
3. Width to Depth ratio of the fire
4. HRR per Unit Area

These parameters are used to produce a two-dimensional probability density function $f_{XY}(x, y)$ that describes the proportion of firebrand particles landing at a given location (x, y) downwind of the fire. This probability density function is described in terms of two component functions:

$$f_{XY}(x, y) = f_X(x)f_{Y|X}(y|x) \quad (7-1)$$

Where $f_X(x)$ describes the marginal probability density function in the downwind x axis, and $f_{Y|X}(y|x)$ describes the marginal probability density function in the crosswind y axis for a given position on the x axis. Each of these functions is described as a composition of existing statistical probability density functions whose parameters depend on the inputs of the model.

7.3 Downwind Marginal Distribution

As described in the previous chapter, the function $f_X(x)$ is described as the sum of a set of probability density functions for each particle type, adjusted by a weighting factor w_p that represents the proportion of particles of each type observed to land over the domain. This function is expressed mathematically as:

$$f_X(x) = \sum_{p=0}^n \frac{w_p}{n} PDF_{Inv.G}(x, \mu_p, \omega_p) \quad (7-2)$$

Where $PDF_{Inv.G}$ is the inverse gaussian probability density function:

$$PDF_{Inv.G}(x, \mu_p, \omega_p) = \frac{1}{\sqrt{2\pi(\frac{x}{\omega_p})^3}} \exp\left(-\frac{\left(\frac{x}{\omega_p} - \mu_p\right)^2}{2\frac{x}{\omega_p}\mu_p}\right) \quad (7-3)$$

The values μ_p and ω_p are determined for each particle type using a linear relationship with the log of the mass of the particle m_p . The slope and intercept of this linear relationship are obtained from a lookup table based on the input parameters – allowing the shape of the function $f_X(x)$ to be adapted to the initial conditions of each scenario. Expressing this relationship mathematically:

$$\mu_p = A * \ln(m_p) + B \quad (7-4)$$

$$\omega_p = C * \ln(m_p) + D \quad (7-5)$$

In Python, these equations can be implemented using publicly available free and open-source function libraries. In particular, the SciPy statistical function library (**scipy.stats**) includes a set of functions for the inverse gaussian distribution. The terms B and D in this set of equations are correlated, lying along a curve that may be approximated as a power law. The coefficients A and C do not show a similarly well-defined relationship.

7.4 Crosswind Conditional Probability Distribution

The crosswind distribution function $f_{Y|X}(y|x)$ is composed of the sum of two skew-normal distributions with parameters determined by simulation initial conditions. These skew-normal distributions are reflections of each other – the shape and location of one distribution has the opposite sign of the other. This is expressed mathematically as:

$$f_{Y|X}(y|x) = \frac{1}{2} PDF_{SN}\left(\frac{y + \varepsilon(x)}{\omega(x)}, \alpha(x)\right) + \frac{1}{2} PDF_{SN}\left(\frac{y - \varepsilon(x)}{\omega(x)}, -\alpha(x)\right) \quad (7-6)$$

The terms α , ε , and ω (shape, location, and scale) are determined by the initial conditions of the scenario. The term α is calculated using a hyperbolic fit based on the forest condition, wind speed, and W:D ratio of the fire. This relationship is expressed as:

$$\alpha = \frac{E_{WD,F}}{x} K_S + H_{W:D,F} \quad (7-7)$$

Where $E_{WD,F}$ and $H_{W:D,F}$ are values based on the W:D ratio of the fire and forest condition, and K_S is a coefficient that adjusts for the effect of wind speed.

The parameter ε determines the position of each half of the distribution along the y -axis. This parameter shows the least consistent trends across the range of cases – so it is difficult to fit a complex function to this parameter. It can be coarsely approximated as corresponding to the position of the edges of the fire at zero meters downwind – but as the distribution moves downwind, the position of the distribution shifts. In forested cases, the ε parameter increases as x increases until it peaks at $\sim 150\text{m}$ downwind of the fire, then it gradually decreases. This contributes to the wider and flatter distribution that we typically observe in forested cases. In open field cases, we observe the opposite trend in ε – it decreases as x increases until it reaches a minimum value at $\sim 150\text{m}$ downwind, then it gradually increases. These trends are approximated using a pair of simple linear functions that describe the position parameter for forested and non-forested cases separately. In the open field case:

$$\varepsilon(x) = \begin{cases} \frac{W}{2} - \frac{R_1 x}{2W} & \text{for } x < 150 \\ \frac{W}{2} + \frac{R_2 x - 150R_1}{2W} & \text{for } x \geq 150 \end{cases} \quad (7-8)$$

In the forested case:

$$\varepsilon(x) = \begin{cases} \frac{W}{2} + \frac{R_3 x}{2W} & \text{for } x < 150 \\ \frac{W}{2} + \frac{(150 - x)R_3}{2W} & \text{for } x \geq 150 \end{cases} \quad (7-9)$$

Where W is the width of the fire in meters and R_n is a constant.

The scale term ω determines the overall size of each half of the distribution in $f_{Y|X}(y|x)$. As the scale term increases, the width of each half increases – with a varying degree of overlap in the centre of the domain. The scale term thus determines both the overall width of the distribution and the magnitude of the peaks of the distribution relative to the centre of the

distribution. The scale term ω is impacted by all scenario parameters to varying degrees – and as such is determined by a more complex function than other parameters. The most significant scenario parameter impacting the scale term is the W:D ratio, with each W:D ratio producing a different trend in ω over downwind distance x . For each W:D ratio, a combination of linear equations were used to describe the trend in ω over x . These equations are used to produce an initial approximation of ω , which is then adjusted for wind speed, forest conditions, and HRR. The initial approximation for ω based on W:D ratio can be expressed as:

$$\omega_{aprx} = 15WD + 0.033x \quad \text{for W:D} = 1 \quad (7-10)$$

$$\omega_{aprx} = \begin{cases} 15WD - 0.3x, & x < 50 \\ 15WD - 15 + 0.13x, & 50 \leq x < 200 \\ 40, & \text{if } 400 \leq x \text{ and Low HRRPUA} \\ 45 + 0.05x, & \text{if } 400 \leq x \text{ and High HRRPUA} \end{cases} \quad \text{for W:D} = 2 \quad (7-11)$$

$$\omega_{aprx} = \begin{cases} 15WD - 0.6x, & x < 50 \\ 15WD - 30 + 0.1x, & 50 \leq x \leq 200 \\ 50, & 200 < x \end{cases} \quad \text{for W:D} = 4 \quad (7-12)$$

This approximation is then multiplied by a correction factor for wind speed:

$$C_{wind} = \frac{0.74U + 26.62}{54} \quad (7-13)$$

In forested cases, a constant is added to ω after applying the wind correction factor to account for the wider and flatter shape of the distribution typically seen in forested conditions.

7.5 Model Performance

7.5.1 Performance Measurement Statistics

Using a combination of qualitative and quantitative statistical tools, we may evaluate the performance of the model described in the previous section and identify its strengths and weaknesses across the range of cases. Similar to the process of fitting the distributions described in the Statistical Methodology chapter, it is necessary to evaluate the performance of the model in a variety of ways so that the location and source of discrepancies between model predictions and FDS numerical simulation results can be identified and corrected.

The most direct qualitative tool for evaluating the performance of the model is to plot the kernel density estimate of a collection of randomly generated points from the model distribution alongside a kernel density estimate of particle landing positions from an FDS simulation. For each set of initial conditions, a large sample of random particle landing positions were generated using the statistical model to produce a dataset similar to that obtained from the FDS simulations. This can be used to produce plots such as Figure 7-1, showing the distribution of particles predicted by the model in blue and the distribution of particles from the FDS simulation in red.




Figure 7-1: Overlaid kernel density estimate plots for the 46km/hr, 2:1 W:D, High HRR, Open Field case.

Using these overlaid KDE plots, we can compare the overall shape of the model output to results obtained from FDS simulations and observe how well the model captures the features of the distribution. In the case shown in Figure 7-1, the model and simulation results are very similar – with the contours near the edge of the fire overlapping and showing nearly identical shapes. The outermost contour of the model follows the shape of the FDS distribution closely, with few areas of either the model or FDS distribution protruding far from each other.

In general, where there are discrepancies between the model and FDS results, it is preferable for the model to over-predict the length and width of the distribution than under-predict. While it is best to minimize the discrepancy, under-prediction of firebrand particle spread is more

likely to lead to negative outcomes in real-world scenarios than over-prediction. Figure 7-2 shows an example of the model over-predicting firebrand spread when compared to an FDS result:




Figure 7-2: Overlaid kernel density estimate plots for the 37km/hr, 1:1 W:D, Low HRR, Open Field case. This case shows an example of the model over-predicting firebrand distribution, especially in the far downwind areas.

In this case, the model has predicted a distribution that projects too many firebrands into the far downwind regions of the domain. The outermost contour of the model stretches well beyond the edges of the FDS distribution, and the second widest model probability contour extends far beyond the second widest contour of the FDS distribution. Only the contours very close to the fire show a similar shape. However, this model output has a reasonable degree of spread in the crosswind direction – the width of the distribution in the y axis is close to correct.

Kernel density estimate plots can also be generated for a single axis, so that the performance of the functions $f_X(x)$ and $f_{Y|X}(y|x)$ can be evaluated separately. Using the same data as the examples shown above, we can generate kernel density estimate plots of the marginal distribution along the x-axis to show the difference in model performance between these two cases. These plots are shown in Figure 7-3.





Figure 7-3: Kernel density estimate plots of the marginal distribution in the x-axis for two cases: (a) the 46km/hr, 4:1 W:D, High HRR, Open Field case. (b) the 37km/hr, 1:1 W:D, Low HRR, Open Field case.

While visual inspection of KDE probability contours and marginal distributions may show obvious discrepancies, to properly evaluate the performance of the model and understand the impact of changes to the functions that determine the shape of the distribution more advanced analytical tools are required. Two additional statistical tools will be used to evaluate the performance of the model along each axis: Quantile-Quantile plots, and a modification of the Cramér-von Mises criterion.

Quantile-Quantile plots (Q-Q plots) were previously used in the statistical methodology chapter, and they are equally applicable here. The Q-Q plot graphs the position of the quantiles of a dataset against the quantiles of another dataset or distribution. If the data is drawn from the same distribution, the quantiles should lie on the diagonal line $y = x$. If the distributions differ, the positions of the quantiles will be shifted above or below the line – the location and direction of the shift in position indicates where and how the distributions differ.

The Cramér-von Mises criterion is a tool for measuring how closely a dataset matches a given distribution (the one-sample Cramér-von Mises test) or comparing two datasets against each other (the two-sample Cramér-von Mises test). These tests depend on calculating the integral of the difference between cumulative distribution functions:

$$\omega^2 = \int_{-\infty}^{\infty} [F_N(x) - F(x)]^2 dF(x) \quad (7-14)$$

The one-sample Cramér-von Mises criterion is the integral of the squared difference between the empirical cumulative distribution function of the dataset and the theoretical distribution, multiplied by the size of the dataset N . The two-sample Cramér-von Mises criterion is the sum of the squared differences between the empirical cumulative distribution functions of each dataset, multiplied by a coefficient $\frac{NM}{N+M}$ where N and M are the size of each dataset[95]:

$$T = \frac{NM}{N+M} \int_{-\infty}^{\infty} [F_N(x) - G_N(x)]^2 dH_{N+M}(x) \quad (7-15)$$

Where $F_N(x)$ and $G_N(x)$ are the empirical cumulative distribution functions of each dataset and $H_{N+M}(x)$ is an empirical cumulative distribution function of both datasets combined. Although we cannot use T to directly compare datasets as each scenario introduces a different

quantity of particles based on the size and HRR of the fire – which impacts the coefficient $\frac{NM}{N+M}$ – we can simply correct for this factor and compare the integral ω^2 between cases.

Model performance over the area of the domain can be quantitatively measured by measuring the Residual Sum of Squares (RSS) between the model prediction and the FDS simulation distribution. This can be approximated by finding the sum of the squared difference between the KDE of the model and the KDE of the simulation result across the domain. This is calculated by evaluating the KDE of the model and the KDE of the simulation across a grid of points that spans the downwind region of the domain¹:

$$RSS = \sum_{x=0}^{800} \sum_{y=-200}^{200} (KDE_{MOD}(x, y) - KDE_{FDS}(x, y))^2 \quad (7-16)$$

The RSS can then be used to calculate a Root Mean Square Error (RMSE) for each case, giving a result that measures the average difference between the model and FDS distributions in the same scale as the distribution[98]. For a grid of n points, the RMSE is calculated as:

$$RMSE = \sqrt{\frac{RSS}{n}} \quad (7-17)$$

While this metric is influenced by the number of points n , the number of points being evaluated is held constant between cases, so RMSE values should be comparable across all cases. The square root of the RSS term can also be used as a measure of the total discrepancy between the model prediction and FDS distribution in the same unit over the whole domain.

We can then produce a table of ω^2 , RSS, and RMSE values to describe the goodness of fit for each case. These goodness of fit measures can also be used to summarize the goodness of fit for groups of cases, to evaluate the performance of the model across different scenario conditions.

¹ The downwind region of the FDS domain spans 400-1200m in the x -axis, the position data in this axis was adjusted by 400m so that it spans 0-800m for simplicity.

7.6 Overall Performance Statistics and Summary

The table of summary statistics describing the goodness of fit for each case individually is Table 7-4, presented below. A set of summary tables describing average goodness of fit metrics for cases grouped by wind speed, W:D ratio, HRR (low/high) and forest condition are presented in Table 7-1 to Table 7-3.

Domain	Speed (km/hr)	Mean ω_x^2	Mean ω_y^2	Mean RSS	$\sqrt{\text{Mean RSS}}$	Mean RMSE %
Open	37	0.01002	0.00192	0.00124	3.53%	0.00607%
	46	0.00789	0.00046	0.00093	3.04%	0.00526%
	54	0.00700	0.00201	0.00081	2.85%	0.00491%
	64	0.00674	0.00199	0.00067	2.59%	0.00453%
Forest	37	0.00739	0.00247	0.00134	3.66%	0.00633%
	46	0.00376	0.00092	0.00123	3.50%	0.00598%
	54	0.00395	0.00144	0.00091	3.02%	0.00519%
	64	0.00432	0.00204	0.00084	2.89%	0.00491%

Table 7-1: Summary statistics of model performance for cases grouped by domain type and wind speed.

Domain	W:D Ratio	Mean ω_x^2	Mean ω_y^2	Mean RSS	$\sqrt{\text{Mean RSS}}$	Mean RMSE %
Open	1	0.01274	0.00194	0.00133	3.65%	0.00636%
	2	0.00176	0.00149	0.00081	2.85%	0.00495%
	4	0.00510	0.00114	0.00057	2.38%	0.00417%
Forest	1	0.00742	0.00307	0.00165	4.06%	0.00713%
	2	0.00134	0.00069	0.00103	3.22%	0.00561%
	4	0.00580	0.00140	0.00055	2.35%	0.00407%

Table 7-2: Summary statistics of model performance for cases grouped by domain type and W:D ratio.

Domain	HRRPUA	Mean ω_x^2	Mean ω_y^2	Mean RSS	$\sqrt{\text{Mean RSS}}$	Mean RMSE %
Open	Low	0.00786	0.00160	0.00111	3.32%	0.00571%
	High	0.00547	0.00151	0.00076	2.76%	0.00479%
Forest	Low	0.00342	0.00191	0.00128	3.57%	0.00615%
	High	0.00629	0.00153	0.00088	2.96%	0.00506%

Table 7-3: Summary statistics of model performance for cases grouped by domain type and HRRPUA level.

From the data presented in these tables, it can be seen that the model has very good performance overall, with small total discrepancies ($\sqrt{\text{Mean RSS}}$) between the model probability distribution function and the probability distribution function produced from FDS simulation data. The

model performs best for cases with higher wind speeds, larger W:D ratios, and higher HRR values. The impact of forested domains on model performance is relatively small, with greater variance in performance between cases in open field domains than between open field and forested cases. These qualities are desirable in addition to good model performance overall – the consistent performance of the model in both open and forested domains and its good performance in more extreme scenarios (larger fires, higher HRRs) may make this model a good basis for further development.

Speed	WD	HRR	Domain	ω_x^2	ω_x	ω_y^2	ω_y	RSS	$\sqrt{\text{RSS}}$ %	RMSE %
37	1	high	open	1.93E-02	13.88%	2.60E-03	5.10%	1.46E-03	3.82%	0.00675%
37	1	low	open	2.24E-02	14.97%	2.59E-03	5.09%	2.30E-03	4.79%	0.00847%
37	2	high	open	2.04E-03	4.52%	4.78E-03	6.92%	6.89E-04	2.62%	0.00464%
37	2	low	open	2.23E-03	4.72%	8.48E-04	2.91%	1.50E-03	3.87%	0.00684%
37	4	high	open	5.55E-03	7.45%	1.74E-04	1.32%	6.30E-04	2.51%	0.00444%
37	4	low	open	8.64E-03	9.29%	5.28E-04	2.30%	8.96E-04	2.99%	0.00529%
46	1	high	open	8.91E-03	9.44%	7.39E-04	2.72%	1.13E-03	3.36%	0.00595%
46	1	low	open	1.58E-02	12.57%	7.83E-04	2.80%	1.66E-03	4.07%	0.00719%
46	2	high	open	1.16E-03	3.41%	2.83E-04	1.68%	7.92E-04	2.81%	0.00497%
46	2	low	open	1.56E-03	3.94%	3.59E-04	1.89%	9.81E-04	3.13%	0.00554%
46	4	high	open	3.10E-03	5.57%	3.24E-04	1.80%	4.57E-04	2.14%	0.00378%
46	4	low	open	4.54E-03	6.74%	2.56E-04	1.60%	5.40E-04	2.32%	0.00411%
54	1	high	open	4.66E-03	6.83%	1.92E-03	4.38%	9.75E-04	3.12%	0.00552%
54	1	low	open	1.40E-02	11.81%	1.93E-03	4.40%	1.47E-03	3.84%	0.00678%
54	2	high	open	1.68E-03	4.10%	7.26E-04	2.69%	6.49E-04	2.55%	0.00450%
54	2	low	open	2.24E-03	4.74%	1.92E-03	4.38%	8.83E-04	2.97%	0.00525%
54	4	high	open	3.58E-03	5.99%	2.81E-03	5.31%	3.98E-04	2.00%	0.00353%
54	4	low	open	5.19E-03	7.21%	2.74E-03	5.23%	4.80E-04	2.19%	0.00387%
64	1	high	open	8.56E-03	9.25%	1.93E-03	4.40%	8.20E-04	2.86%	0.00506%
64	1	low	open	8.37E-03	9.15%	3.00E-03	5.47%	8.63E-04	2.94%	0.00519%
64	2	high	open	1.64E-03	4.05%	3.47E-04	1.86%	4.02E-04	2.01%	0.00355%
64	2	low	open	1.49E-03	3.86%	2.67E-03	5.17%	5.96E-04	2.44%	0.00431%
37	1	high	forest	8.77E-03	9.37%	9.33E-03	9.66%	1.41E-03	3.75%	0.00663%
37	1	low	forest	1.45E-02	12.06%	3.73E-03	6.11%	2.21E-03	4.70%	0.00830%
37	2	high	forest	1.35E-03	3.68%	7.68E-04	2.77%	1.26E-03	3.54%	0.00627%
37	2	low	forest	7.56E-04	2.75%	3.91E-04	1.98%	1.63E-03	4.03%	0.00713%
37	4	high	forest	1.83E-02	13.54%	1.88E-04	1.37%	5.23E-04	2.29%	0.00404%
37	4	low	forest	6.14E-04	2.48%	3.95E-04	1.99%	1.00E-03	3.17%	0.00560%
46	1	high	forest	4.01E-03	6.33%	1.53E-03	3.91%	1.74E-03	4.17%	0.00737%
46	1	low	forest	6.04E-03	7.77%	2.34E-03	4.83%	2.27E-03	4.76%	0.00842%
46	2	high	forest	1.07E-03	3.27%	3.23E-04	1.80%	8.08E-04	2.84%	0.00502%
46	2	low	forest	6.42E-04	2.53%	4.40E-04	2.10%	1.32E-03	3.63%	0.00641%
46	4	high	forest	9.55E-03	9.77%	1.10E-04	1.05%	4.22E-04	2.06%	0.00363%
46	4	low	forest	1.23E-03	3.51%	7.88E-04	2.81%	8.09E-04	2.84%	0.00503%
54	1	high	forest	6.06E-03	7.79%	2.01E-03	4.48%	1.21E-03	3.47%	0.00614%
54	1	low	forest	5.53E-03	7.44%	2.57E-03	5.07%	1.52E-03	3.90%	0.00689%
54	2	high	forest	1.12E-03	3.34%	4.25E-04	2.06%	6.74E-04	2.60%	0.00459%
54	2	low	forest	2.10E-03	4.59%	1.17E-03	3.42%	1.11E-03	3.33%	0.00589%
54	4	high	forest	7.22E-03	8.49%	3.72E-04	1.93%	3.43E-04	1.85%	0.00327%
54	4	low	forest	1.68E-03	4.09%	2.07E-03	4.55%	6.10E-04	2.47%	0.00437%
64	1	high	forest	9.80E-03	9.90%	1.21E-03	3.48%	1.24E-03	3.53%	0.00623%
64	1	low	forest	4.62E-03	6.80%	1.81E-03	4.26%	1.59E-03	3.98%	0.00704%
64	2	high	forest	1.88E-03	4.34%	6.48E-04	2.54%	6.05E-04	2.46%	0.00435%
64	2	low	forest	1.83E-03	4.28%	1.32E-03	3.63%	8.73E-04	2.96%	0.00522%
64	4	high	forest	6.33E-03	7.96%	1.39E-03	3.73%	3.14E-04	1.77%	0.00313%
64	4	low	forest	1.46E-03	3.82%	5.87E-03	7.66%	3.95E-04	1.99%	0.00351%

Table 7-4: *Quality of fit for each case, determined by the magnitude of ω^2 for each axis. The square roof of ω^2 is presented as a percentage.*

The measures of goodness of fit in the x and y axes (ω^2) show a higher-than-average discrepancy in the x axis marginal distribution for cases with a small W:D ratio, especially in open field cases. Discrepancies in the y axis marginal distribution are highest for cases with small W:D ratios in forested domains. These small W:D cases also have the poorest overall quality of fit of all cases in general, as measured by $\sqrt{\text{RSS}}$. The lower performance of the model here could be due to the physics of particle transport in these cases. The trends in particle landing position that the model is attempting to capture are the result of physical processes driven by the convective plume. In the cases with small fires and lower HRRs per unit area, the impact of these physical processes may be weaker than the model predicts, leading to a tendency to over-predict dispersion in the x -axis. The under-prediction of spread in the y -axis for some narrow fires may be due to an underestimation of the impact of open domains on the location term in $f_{Y|X}(y|x)$.

The best quality of fit by $\sqrt{\text{RSS}}$ is observed for the 2:1 and 4:1 W:D cases with high HRRs in both open field and forested domains. From a physics perspective, this is likely because these cases produce consistently larger and more powerful convective plumes and vortices. The more powerful plume and vortices produce more consistent effects on particle transport that are easier to represent statistically. In addition to this, these cases often had larger datasets generated from FDS simulations – particle insertion rates were determined by the size and HRR of the fire, and the larger and more powerful plume may have lofted more particles downwind of the burning region to be captured in the distribution. The larger size of the particle landing dataset may have improved the quality of fit for these cases.

While these summary statistics may show that the model has good performance overall and provide a means to identify the best- and worst-case scenarios for the model, the nuances of how the model performs for each case can only be shown by examining the output of each case individually. For this purpose, a selection of model results representing a range of model performance metrics in forested and open field conditions will be presented with commentary on the quality of fit. The remainder of the results will be presented in an appendix.

7.7 Model Performance in Specific Scenarios

The scenarios shown in this section are intended to represent a wide range of model performance across all conditions, showing an example of the best, worst, and average case performance. This is to indicate the range of performance that can be expected from the model and to highlight where improvements could be made. As shown in the statistical summary tables, the average quality of the model output is generally very good. The detailed analysis of these scenarios will also illustrate how the summary statistics presented in Table 7-4 can be interpreted to properly understand how well the model and simulation distributions agree.

7.7.1 Open Field Cases

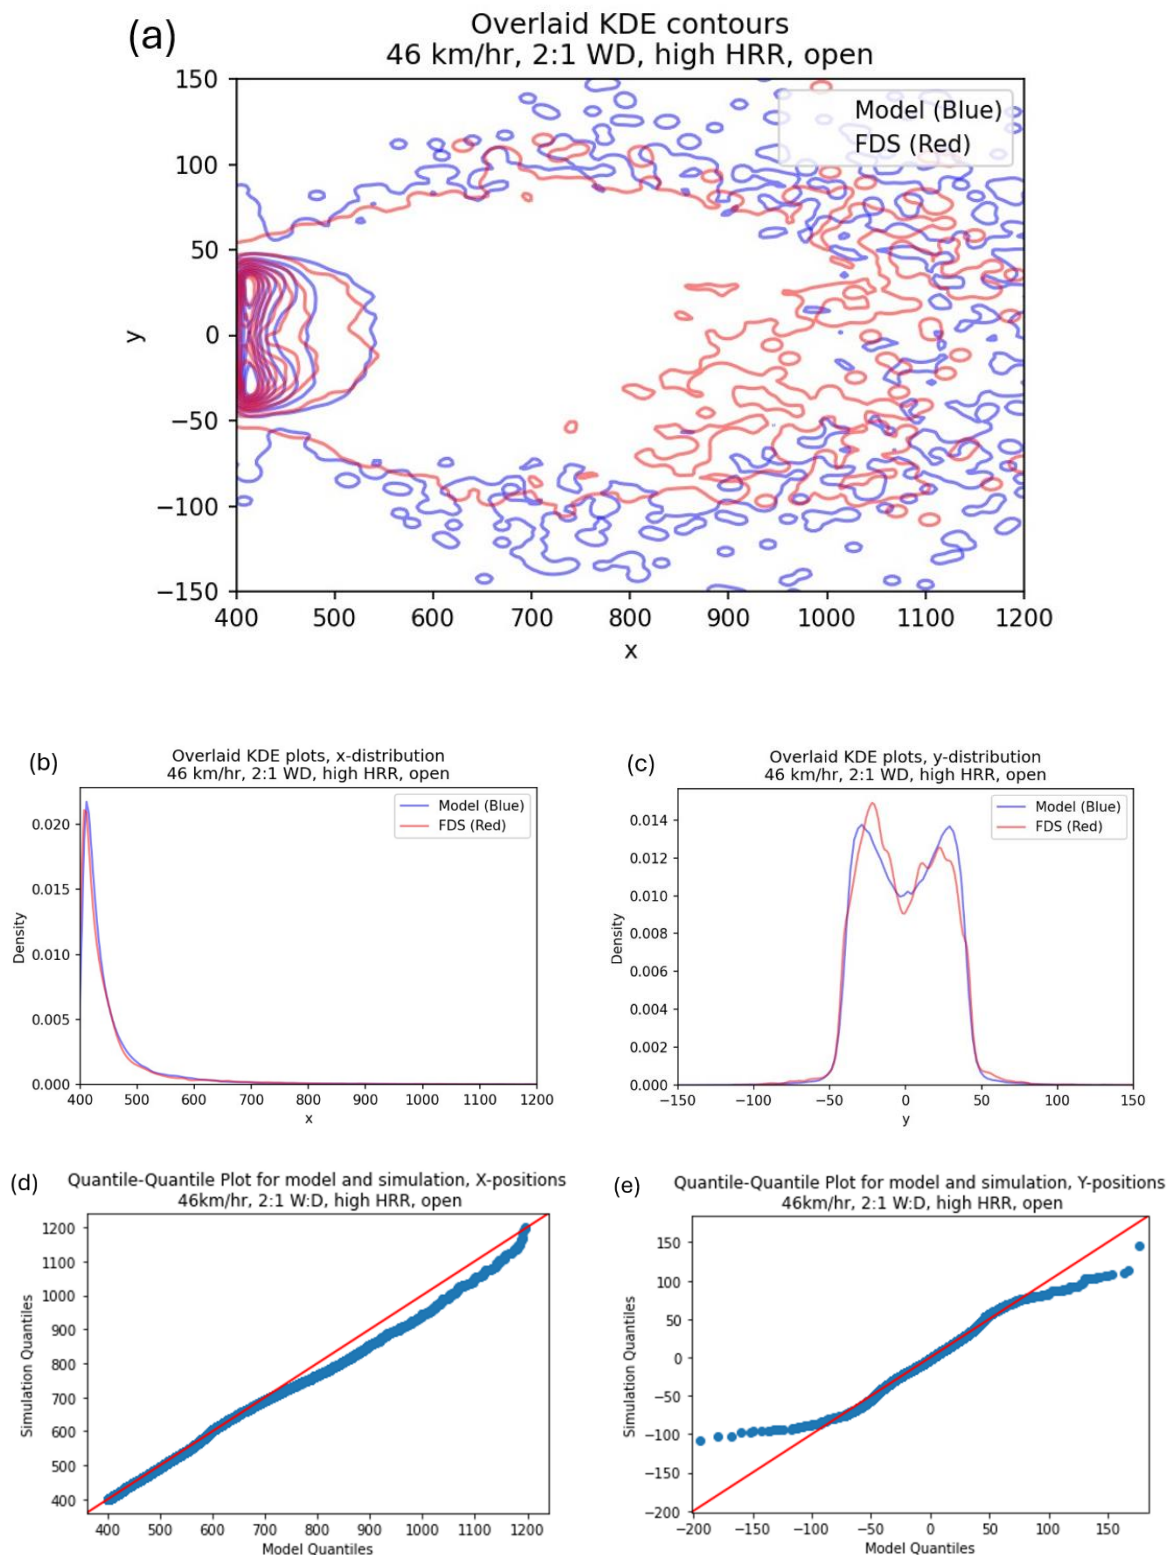


Figure 7-4: Plots of model predictions for the 46km/hr, 2:1 W:D, High HRR, Open Field case. (a) Overlaid KDE for the downwind region of the domain. (b) Overlaid KDE for the marginal distribution in the x-axis. (c) Overlaid KDE for the marginal distribution in the y-axis. (d) Q-Q plot comparing the model and simulated distribution in the x-axis (e) Q-Q plot comparing the model and simulated distribution in the y-axis.

Key Statistics: $\sqrt{RSS} = 2.81\%$, $\omega_x^2 = 1.16E-03$, $\omega_y^2 = 2.83E-04$

The 46km/hr, 2:1 W:D, High HRR, Open Field case is a useful example of how the summary statistics in Table 7-1 can be interpreted to understand the quality of fit produced by the model. While this case has close to average \sqrt{RSS} for the set of open field cases, this case has some of the most accurate model results in both marginal distributions (ω_x^2 and ω_y^2). This produces a result that fully captures the extent of the distribution in both axes. The distribution obtained from FDS falls almost entirely within the model prediction, and the probability density contours of the model align closely with the contours of the FDS distribution. Similar model performance is achieved in a number of other open field cases with comparably low \sqrt{RSS} , ω_x^2 and ω_y^2 values, such as those described in Table 7-5:

Speed	W:D	HRR	Domain	ω_x^2	ω_x	ω_y^2	ω_y	$\sqrt{RSS}\%$	RMSE%
37	4	high	open	5.55E-03	7.45%	1.74E-04	1.32%	2.51%	0.00444%
46	2	low	open	1.56E-03	3.94%	3.59E-04	1.89%	3.13%	0.00554%
46	4	high	open	3.10E-03	5.57%	3.24E-04	1.80%	2.14%	0.00378%
46	4	low	open	4.54E-03	6.74%	2.56E-04	1.60%	2.32%	0.00411%
64	2	high	open	1.64E-03	4.05%	3.47E-04	1.86%	2.01%	0.00355%

Table 7-5: Scenarios with similar quality of fit to the 46km/hr, 2:1, High HRR, Open Field case

These cases show similar performance to the case shown in Figure 7-4 as measured by the statistical goodness of fit metrics – in particular, the product of ω_x^2 and ω_y^2 in these cases are close to those of the above case.

In some of these cases, the FDS simulated distribution is slightly skewed towards one side of the y-axis in the far downwind regions of the domain, likely due to transient wake vortices shifting the distribution of particles as they move through the domain. This slightly increases the \sqrt{RSS} and ω_y^2 metrics for these cases. However, as this skew primarily occurs in regions with low particle landing density, the effect is quite small.

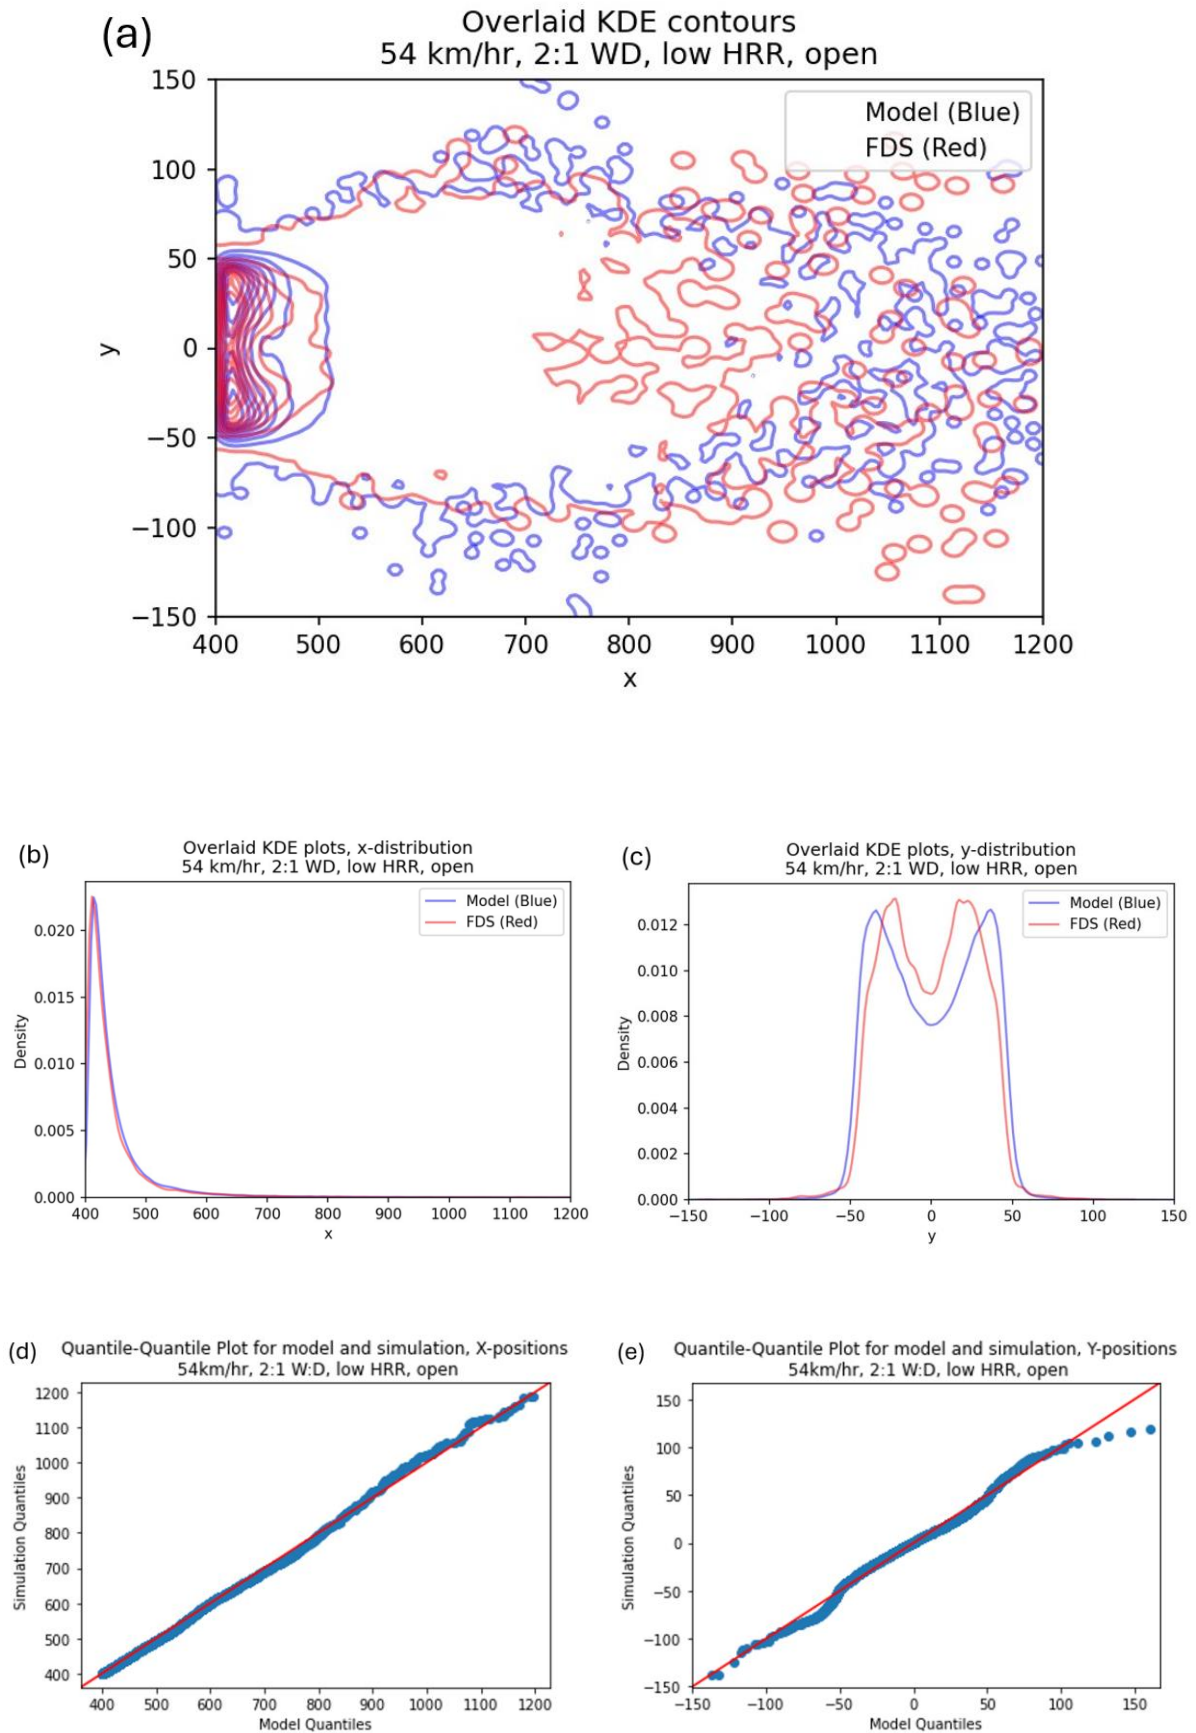


Figure 7-5: Plots of model predictions for the 54km/hr, 2:1 W:D, Low HRR, Open Field case. (a) Overlaid KDE for the downwind region of the domain. (b) Overlaid KDE for the marginal

distribution in the x-axis. (c) Overlaid KDE for the marginal distribution in the y-axis. (d) Q-Q plot comparing the model and simulated distribution in the x-axis (e) Q-Q plot comparing the model and simulated distribution in the y-axis.

Key Statistics: $\sqrt{RSS} = 2.97\%$, $\omega_x^2 = 2.24E-03$, $\omega_y^2 = 1.92E-03$

This case (54km/hr, 2:1 W:D, Low HRR, Open Field) has a \sqrt{RSS} that is very close to the average for open field cases and a product of ω_x^2 and ω_y^2 that is close to the median of open field cases. While the model distribution covers the entire span of the simulated distribution, we can see from the KDE plot of the marginal distribution in the y-axis (Figure 7-5c) that the model distribution slightly over-predicts particle landing density near the edges of the distribution, and slightly under-predicts the landing density in the centre. This type of discrepancy is typical of the average performance of the model. However, the discrepancy in this case is a slight over-prediction of spread – leading to a more conservative prediction of what areas may be at risk of firebrand attack. This is preferable to the alternative of under-prediction. Other cases similar in performance to this case are shown in Table 7-6:

Speed	W:D	HRR	Domain	ω_x^2	ω_x	ω_y^2	ω_y	$\sqrt{RSS}\%$	RMSE%
37	2	high	open	2.04E-03	4.52%	4.78E-03	6.92%	2.62%	0.00464%
37	2	low	open	2.23E-03	4.72%	8.48E-04	2.91%	3.87%	0.00684%
37	4	low	open	8.64E-03	9.29%	5.28E-04	2.30%	2.99%	0.00529%
46	1	high	open	8.91E-03	9.44%	7.39E-04	2.72%	3.36%	0.00595%
54	1	high	open	4.66E-03	6.83%	1.92E-03	4.38%	3.12%	0.00552%
54	2	high	open	1.68E-03	4.10%	7.26E-04	2.69%	2.55%	0.00450%
54	4	high	open	3.58E-03	5.99%	2.81E-03	5.31%	2.00%	0.00353%
54	4	low	open	5.19E-03	7.21%	2.74E-03	5.23%	2.19%	0.00387%
64	2	low	open	1.49E-03	3.86%	2.67E-03	5.17%	2.44%	0.00431%

Table 7-6: Scenarios with similar quality of fit to the 54km/hr, 2:1 W:D, Low HRR, Open Field case.

These cases can be considered representative of the median performance of the model. In general, a good quality of fit as measured by $\sqrt{RSS}\%$ is achieved, but the product of ω_x^2 and ω_y^2 is larger than in the best-case scenarios. This means that the model distribution is overall a good match for the simulated distribution, but there are some discrepancies that can be identified along the marginal distributions. These discrepancies are usually larger along the x-axis, particularly in cases with narrow W:D ratios. This may indicate that the functions describing particle distribution in the x-axis could be improved in the median case.

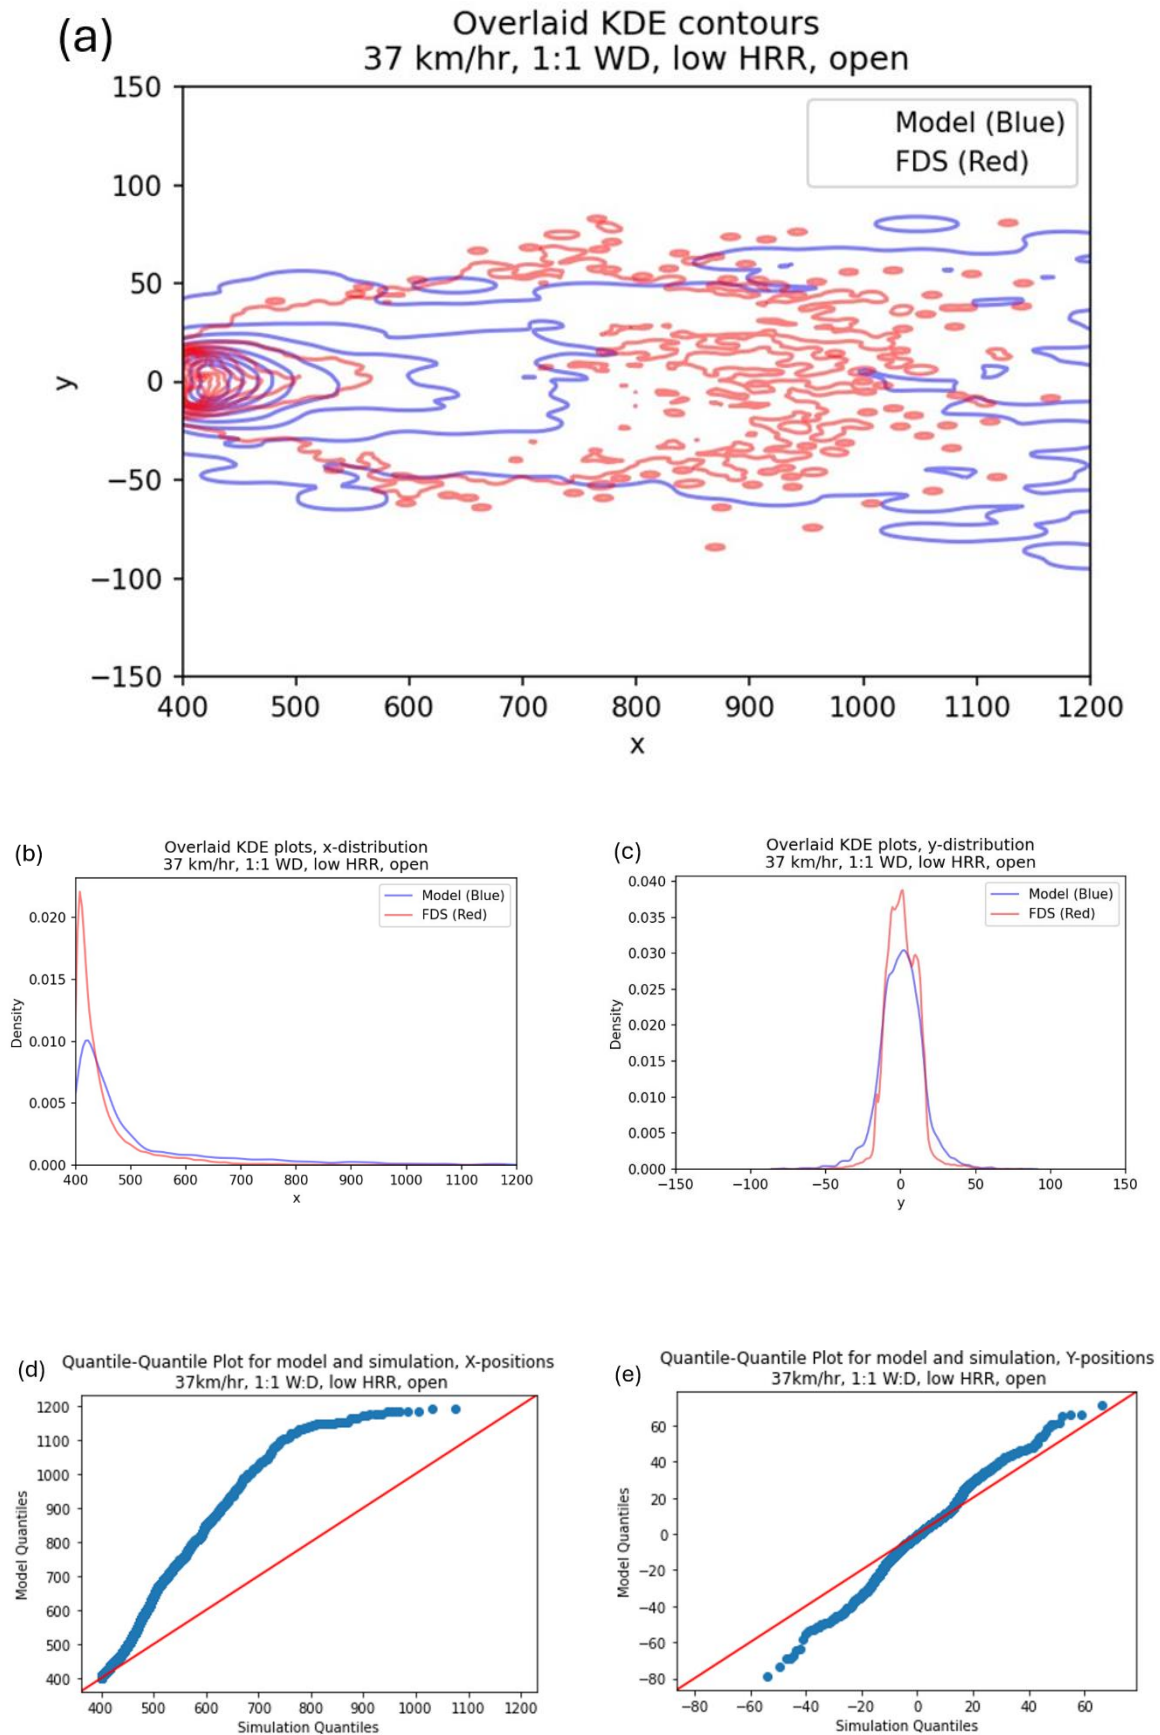


Figure 7-6: Plots of model predictions for the 37km/hr, 1:1 W:D, Low HRR, Open Field case. (a) Overlaid KDE for the downwind region of the domain. (b) Overlaid KDE for the marginal

distribution in the x-axis. (c) Overlaid KDE for the marginal distribution in the y-axis. (d) Q-Q plot comparing the model and simulated distribution in the x-axis (e) Q-Q plot comparing the model and simulated distribution in the y-axis.

Key Statistics: $\sqrt{RSS} = 4.79\%$, $\omega_x^2 = 2.24E-02$, $\omega_y^2 = 2.59E-03$

This case (37km/hr, 1:1 W:D, Low HRR) has the highest \sqrt{RSS} value of all open field cases – more than two standard deviations above the mean. This model result over-predicts particle dispersion in the x-axis considerably, leading to an outermost probability density contour that extends well beyond the simulated distribution we are trying to model. Only one other open field case has a similar \sqrt{RSS} – the 37km/hr, 1:1 W:D, High HRR, Open Field case, with nearly identical initial conditions.

Other small and low-intensity fires at higher wind speeds (e.g. 1:1 W:D, Low HRR cases at 46km/hr and 54km/hr) also have larger than average \sqrt{RSS} values, but not to the same degree as this case. These two cases are also unusual for their under-prediction of spread in the y-axis. This indicates that while model performance is good in the average case and high intensity fire cases, performance could be improved for particularly small and low-intensity fires in open domains.

7.7.2 Forested Cases








Figure 7-7: Plots of model predictions for the 46km/hr, 1:1 W:D, High HRR, Forested case. (a) Overlaid KDE for the downwind region of the domain. (b) Overlaid KDE for the marginal distribution in the x-axis. (c) Overlaid KDE for the marginal distribution in the y-axis. (d) Q-Q plot comparing the model and simulated distribution in the x-axis (e) Q-Q plot comparing the model and simulated distribution in the y-axis.

Key Statistics: $\sqrt{RSS} = 4.17\%$, $\omega_x^2 = 4.01E-03$, $\omega_y^2 = 1.53E-03$

The 46km/hr, 1:1 W:D, High HRR, Forested case provides a useful example of a 1:1 W:D case with better model performance than the open field case shown in Figure 7-6. Although this case also has a higher than average \sqrt{RSS} for the set of forested cases, the ω_x^2 and ω_y^2 terms are greatly improved when compared to the open field case shown in Figure 7-6. This is especially clear in the overlaid KDE plots and Q-Q plot shown in Figure 7-7b and Figure 7-7d – the distribution of particles in the x -axis in this case is dramatically improved.

While this improvement in ω_x^2 relative to open field cases is seen for most 1:1 W:D cases in forested domains, narrow fire scenarios with low HRRs continue to have a lower than average quality of fit when compared to wider fire cases.

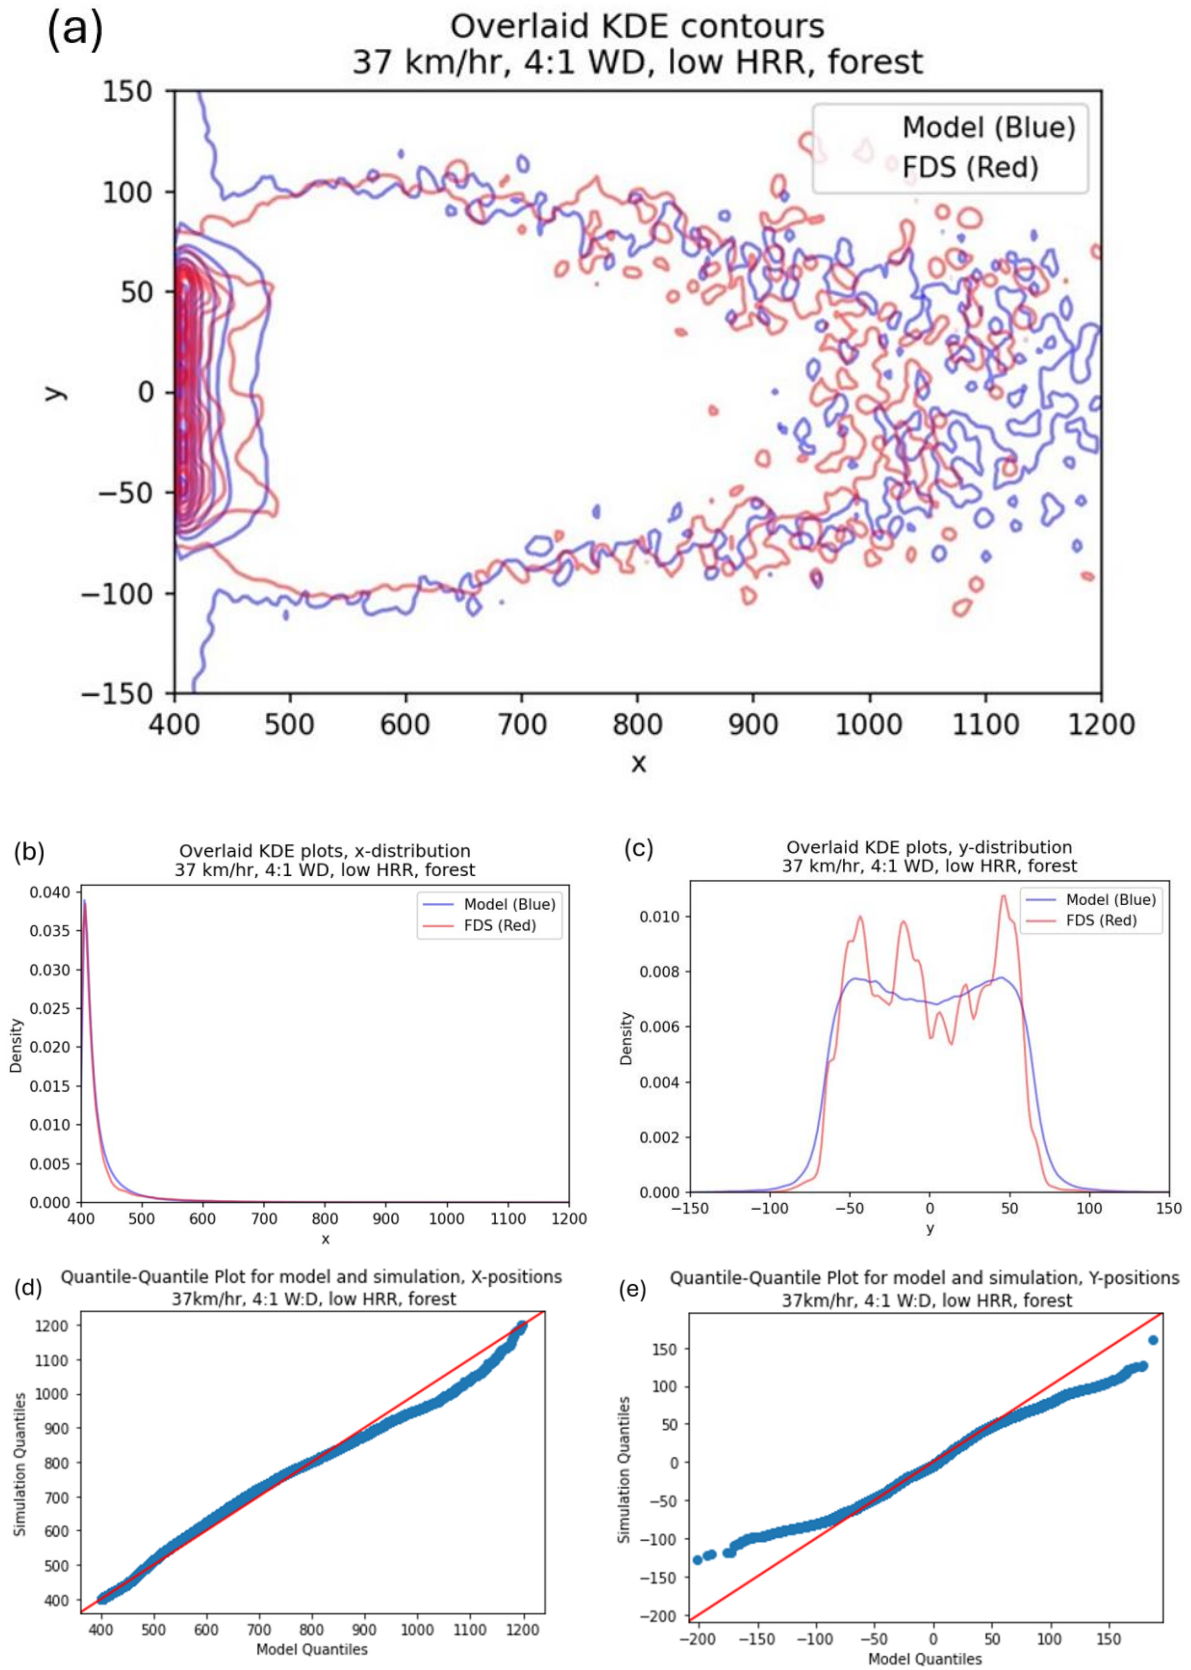


Figure 7-8: Plots of model predictions for the 37km/hr; 4:1 W:D, Low HRR, Forested case. (a) Overlaid KDE for the downwind region of the domain. (b) Overlaid KDE for the marginal distribution in the x-axis. (c) Overlaid KDE for the marginal distribution in the y-axis. (d) Q-

Q plot comparing the model and simulated distribution in the x-axis (e) Q-Q plot comparing the model and simulated distribution in the y-axis.

Key Statistics: $\sqrt{RSS} = 3.17\%$, $\omega_x^2 = 6.14E-04$, $\omega_y^2 = 3.95E-04$

This case (37km/hr, 4:1 W:D, Low HRR, Forested) is another example of a case with very good model performance in the marginal distributions, but average performance in the \sqrt{RSS} statistic. This model prediction has some of the lowest ω_x^2 and ω_y^2 values of any case in either domain type, and the contours of the model follow the shape of the simulated distribution extremely closely. The \sqrt{RSS} value may be higher than the quality of the fit would suggest due to sudden changes in the shape of the FDS marginal distribution, seen in Figure 7-8c. These sharp fluctuations in particle landing density are effectively a source of noise that the model cannot predict, increasing the average difference between the model PDF and the PDF generated from FDS simulation data.

In addition to this source of noise, the model has introduced a small population of particles distributed along the leftmost edge of the domain that are not present in the FDS simulation. This scattering of particles along the leftmost edge of the domain is caused by a high value of the scale parameter of $f_{Y|X}(y|x)$ and high particle landing density in the region close to the fire. As the function $f_{Y|X}(y|x)$ is composed of a sum of two skew-normal distributions with opposite skewness and location, distributions with a wide and flat marginal distribution in the y-axis are prone to this source of discrepancy between the model and FDS simulation results. This issue could be resolved by using a truncated skew-normal distribution as the basis for $f_{Y|X}(y|x)$ in the region close to the fire. Future iterations of this model may test this alternative composition of $f_{Y|X}(y|x)$.

While this case has a particularly good quality of fit, there are several other cases in the forested domain with similarly good performance in the quality of the marginal distributions, described in Table 7-7:

Speed	W:D	HRR	Domain	ω_x^2	ω_x	ω_y^2	ω_y	$\sqrt{RSS}\%$	RMSE%
37	2	low	forest	7.56E-04	2.75%	3.91E-04	1.98%	4.03%	0.00713%
46	2	low	forest	6.42E-04	2.53%	4.40E-04	2.10%	3.63%	0.00641%
46	4	low	forest	1.23E-03	3.51%	7.88E-04	2.81%	2.84%	0.00503%
54	2	high	forest	1.12E-03	3.34%	4.25E-04	2.06%	2.60%	0.00459%
64	2	high	forest	1.88E-03	4.34%	6.48E-04	2.54%	2.46%	0.00435%

Table 7-7: Scenarios with similar quality of fit to the 37km/hr, 4:1 W:D, Low HRR, Forested scenario.

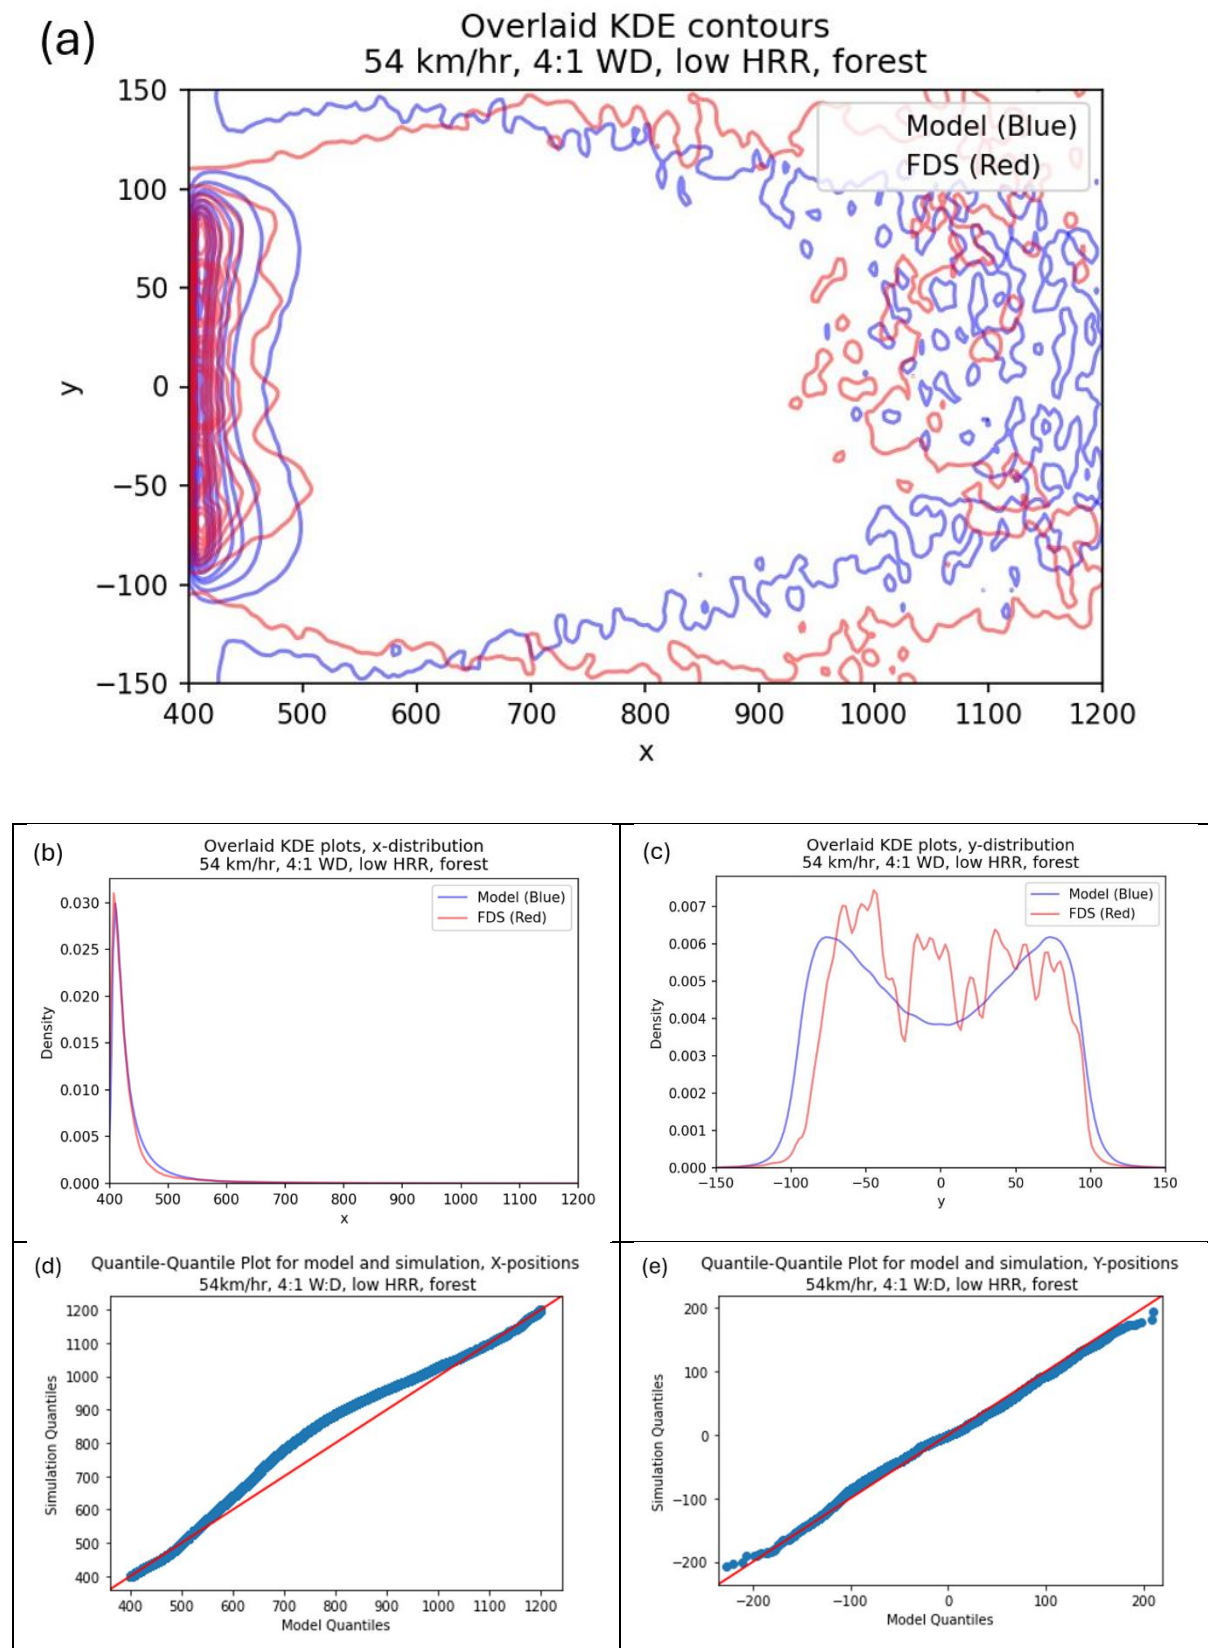


Figure 7-9: Plots of model predictions for the 54km/hr, 4:1 W:D, Low HRR, Forested case. (a) Overlaid KDE for the downwind region of the domain. (b) Overlaid KDE for the marginal

distribution in the x-axis. (c) Overlaid KDE for the marginal distribution in the y-axis. (d) Q-Q plot comparing the model and simulated distribution in the x-axis (e) Q-Q plot comparing the model and simulated distribution in the y-axis.

Key Statistics: $\sqrt{RSS} = 2.47\%$, $\omega_x^2 = 1.68E-03$, $\omega_y^2 = 2.07E-03$

This case is an example of a fairly typical modelling result in the forested domain. The \sqrt{RSS} and ω_x^2 and ω_y^2 values in this case are slightly lower than the average for forested cases, suggesting a high quality of fit. The quality of fit in the marginal distribution along the y-axis is very good, as shown in the Q-Q plot (Figure 7-9e). However, while the magnitude of the difference between the model probability distribution function and the distribution of particles obtained from FDS simulation is small, this case shows a small under-prediction of particle dispersion at long ranges along the y-axis.

While there are a few other instances of under-prediction of particle dispersion across the range of scenarios, the majority of cases tend to either over-predict particle dispersion or closely follow the shape of the FDS simulation data. Future modifications of this model may aim to eliminate or minimize under-prediction errors by iterative refinement.

7.8 Conclusion

For the range of scenarios we have chosen to run, the model performs very well overall. The average \sqrt{RSS} for open field cases is 2.97% with a standard deviation of 0.71%, indicating that the model PDF closely follows the PDF produced from simulation data in the average case. Performance in forested domains is similar to open field cases, with similar average \sqrt{RSS} (Avg. 3.17%, S.D. 0.85%) and ω^2 . These results show that the model is able to reproduce the characteristics of the FDS simulated distributions with a high degree of fidelity while using a small fraction of the computational resources of a CFD simulation - the execution time of the algorithm that produces the model dataset is on the order of 1-3 minutes using a single core of a commercial desktop CPU (AMD Ryzen 5800X3D). This time could be improved dramatically with further refinement of the model algorithm or the use of multiple CPU threads. To illustrate the difference in computation time required, the CPU time required to complete an FDS simulation of particle landing positions and the CPU time required to run the model is presented for an example case below:

Speed	W:D	HRR	Domain	Particles	FDS Real	FDS CPU	Model Real	Model CPU
54	1	low	forest	142,360	19:11:40	1842:40:00	00:01:31	00:01:31

Table 7-8: Time requirements for FDS simulation and model output for an example case. The FDS simulation required 19h 11m of real time with 96 CPU cores at the National Computing Infrastructure Gadi supercomputer, using Intel Xeon Cascade Lake CPUs. The model required 1m 31s on a single CPU core of a desktop PC, using an AMD Ryzen 5800X3D.

There are a small number of cases where the model over-predicts or under-predicts particle dispersion in the downwind direction – these are almost all cases with 1:1 W:D ratio in the open field domains. This suggests that the function $f_X(x)$ could be improved in these cases – either by fitting a new set of parameters for these cases, or by modifying the function itself. These cases produce the largest ω^2 values in the x distribution and the discrepancy is clearly visible on Q-Q plots for these datasets. In some of these cases, there is also a tendency to under-predict the dispersion of particles in the crosswind direction at intermediate distances downwind of the fire. This is likely due to the location function in $f_{Y|X}(y|x)$ under-predicting the tendency for the position of the halves of the distribution to shift, or due to the skewness function over-estimating skewness in these cases. Other opportunities to improve the quality of the model have been identified, in particular for cases with high W:D and flat marginal distributions in the y -axis. The use of a truncated skew-normal distribution in the function $f_{Y|X}(y|x)$ for regions close to the fire may eliminate some sources of error in the model in these cases.

Chapter 8 Summary and Conclusion

Contents

Summary and Conclusion	212
Overview of Model Development	213
Physics-based Model Selection and Scenario Development	213
Simulation Results and Analysis of Physics	213
Statistical Methodology and Model Development	214
Model Performance.....	216
Recommendations for Future Research	217

8.1 Introduction

The primary aim of this research was to develop a statistical model of firebrand landing based on results obtained from physics-based modelling of a realistic wildfire scenario. A rigorous study was conducted to develop this model, investigating the problem of modelling the distribution of firebrand particles from both a physical and statistical perspective. The model developed from this research shows good performance characteristics – producing particle distributions that closely match results obtained from physics-based simulation. The model requires few input parameters, and the trends identified in the shape of the distribution of particles may be generalizable to a broader range of scenarios.

The development of this simplified model of firebrand landing represents a significant contribution to knowledge within the field of firebrand modelling. This contribution to knowledge includes both the model itself and key findings about the behaviour of firebrands obtained from physics-based simulation results. The distribution of firebrands over short ranges has been shown to be well-approximated by a set of relatively simple functions whose parameters can be estimated based on a small number of easily measured physical conditions. The underlying physics that causes these patterns of firebrand distribution to occur has been investigated, primarily relating to the motion of firebrand particles as they descend through a counter-rotating vortex pair beneath the convective plume.

8.2 Overview of Model Development

8.2.1 Physics-based Model Selection and Scenario Development

The model development process began with the selection of an experimentally validated physics-based model for use as a basis for modelling firebrand particle transport. While many physics-based models may be capable of simulating the airflow around a convective plume in a crosswind and modelling the transport of particles from a simulated fire, recent work by Wadhwani [22] experimentally validating FDS with the Haider and Levenspiel drag model against firebrand distribution from a firebrand dragon made it an ideal choice. In addition to this, FDS includes a more detailed model of combustion and the interaction of chemical species in a fire scenario than many other physics-based modelling packages.

Having chosen this model, a range of wildfire scenarios were designed, tested, and simulated to produce a dataset of firebrand particle trajectories and landing positions for different initial wildfire conditions, and to produce data on the physical conditions that produced these distributions. These scenarios were designed to represent a range of McArthur Mk.5 Forest Fire Danger Index (FFDI) values [49, 63], ranging from “Severe” to “Extreme” fire conditions. These FFDI values determine a range of conditions that are held constant between scenarios (e.g. ambient temperature, humidity) as well as the wind speed at the inlet of the domain.

With the wind speed determined by the set FFDI values, the dimensions of the fire were chosen to represent a range of different sizes. The flame depth is calculated by assuming a constant residence time and forward rate of spread. The rate of spread is calculated based on the Project Vesta[36] rate of spread equations for a eucalyptus forest with surface and near-surface fuel conditions similar to a forest with a fuel age of ~240 months (as described in Project Vetsa: Fire in Dry Eucalypt Forest, Appendix II). The final dimensions of the fire are then determined by assigning a range of width to depth ratios (W:D ratio) to each scenario. Finally, the HRR of the fire is specified, based on two HRRs Per Unit Area (HRRPUA) – a high HRRPUA of 1361 kW/m², and a low HRRPUA of 899 kW/m². Each of these conditions (wind speed, flame dimensions, HRR) were then applied to two domain types: open field and forested.

8.2.2 Simulation Results and Analysis of Physics

Data on airflow and particle motion in each of the simulation scenarios was analyzed to determine the underlying physical mechanisms that led to the observed distribution of particles. Analysis of the shape of the convective plume showed that scenarios in forested domains tended to produce convective plumes that closely followed theoretical models of plume rise,

while the plume in open field scenarios tended to lie beneath theoretical model predictions. This may be caused by airflow in open field scenarios interacting with the edge of the forest, similar to airflow over a step.

Analysis of airflow around and beneath the convective plume in these scenarios indicated the formation of large vortex structures beneath the plume, a finding consistent with empirical observations from real wildfires and wind tunnel experiments. These vortices are identified as consistently influencing the trajectory of firebrand particles to form predictable firebrand landing distributions.

In particular, the Counter-rotating Vortex Pair (CVP) that formed beneath the convective plume was found to have a particularly significant impact on the trajectory of particles. The position and velocity components of airflow in the CVP influenced the trajectory of particles in a predictable way, especially in open field scenarios – a distinct relationship between particle initial position and landing position was identified. The influence of the CVP was especially pronounced in the trajectory of particles near the extreme edges of the fire in the y -axis. Particles in these areas showed a tendency to initially travel far outwards from the fire, as they interacted with the upper regions of the CVP. As these particles lost altitude, they passed through the lower regions of the CVP and moved inwards, towards the center of the domain. This outwards-then-inwards trajectory was observed in all cases.

The position and velocity components of the CVP depend on the characteristics of the fire and the conditions of the domain – with the presence or absence of a forest in the region downwind of the fire producing a particularly large change in the position and velocity of airflow in the vortices. Understanding how scenario conditions impact vortex behaviour – and thus particle trajectory – provides a physical basis for interpreting observed statistical trends in particle distribution.

8.2.3 Statistical Methodology and Model Development

Particle landing positions were recorded as a collection of points distributed over the area downwind of the fire in each scenario. These distributions were then analyzed statistically to identify patterns in the density of particles landing over the domain. One of the most significant observations was that many scenarios produced a distinct two-peaked pattern of particle landing density. In many open field cases, especially those with larger $W:D$ ratios, particles were found to land in a two-peaked shape with high landing densities on either side of the center line of the domain. Some experimental observations of two-peaked distributions of

short-range firebrands were made in Project Vesta[36], but these shapes were not described statistically.

In addition to this, the distribution of firebrands along the downwind x -axis was observed to follow a shape generally similar to some models reported in literature – a rapid increase in density, peaking a short distance from the edge of the fire, followed by a long gradual decrease in particle landing density. A range of models described in literature were tested against the dataset of particle landing positions in the x -axis. Of the distributions of firebrand landing described in literature, the distribution of particles over the x -axis was most similar to a log-normal function – but the quality of fit could be improved.

To describe the distribution of particles in two dimensions, it was necessary to find a pair of functions that would optimally describe the distribution of particles over the x -axis and the distribution of particles over the y -axis for a given position on the x -axis. These functions are $f_X(x)$ and $f_{Y|X}(y|x)$, respectively.

For the function describing the distribution over the x -axis $f_X(x)$, a range of candidate functions from literature and similar functions were selected and tested. Each candidate function for the x -axis distribution was fit to the data using a Maximum Likelihood Estimation (MLE) approach and evaluated using a range of qualitative and quantitative statistical tools to evaluate goodness of fit. Two functions were found with near-identical goodness of fit – the Inverse Gaussian distribution and log-T distribution. The Inverse Gaussian function was selected for further development based on inspection of Q-Q plots and simpler mathematical properties.

For the distribution of particles in the y -axis, a function that could smoothly transition across the variety of distribution shapes observed across the range of scenarios was required. Many cases had produced clearly single-peaked distributions in the y -axis, while other cases were clearly two-peaked. Some cases, such as those with larger W:D ratios in forested domains, produced wide and relatively flat distributions in the y -axis. Given this requirement, it was unlikely that any well-known closed-form probability density function would adequately fit the data, and a new function would need to be constructed.

The process of developing a function for $f_{Y|X}(y|x)$ began by investigating a function for $f_Y(y)$ – a function that describes the distribution of particles in the y -axis over the whole domain, rather than for a particular position in the x -axis. An initial approximation for this function was a gaussian mixture – the sum of two partly overlapping normal distributions. By changing the

value of μ and σ for each component of this distribution, it could smoothly transition between single and double peaked. However, the tails of the gaussian mixture distribution were a poor match for the simulation data, and the locations of the peaks in the distribution rarely matched closely. The gaussian mixture function was also a particularly poor fit for the wide and flat distributions observed in high W:D forested cases.

Through an iterative process, a greatly improved quality of fit was obtained for the entire range of scenarios by using a mixture of Skew-normal distributions rather than gaussian distributions. This function could smoothly transition between single and double peaked shapes, and could produce the wide and flat distributions observed in high W:D forested scenarios. As the initial conditions of the scenarios did not contain any terms that would produce a consistent asymmetry in the distribution, the Skew-normal mixture distribution was modified so that the parameters of each component were equal and opposite to one another – each component of the distribution would be a reflection of the other over the line $y = 0$. This produced a perfectly symmetrical distribution that performed very similarly to the asymmetrical skew-normal mixture distribution, and halved the number of terms to be fit for each scenario. With a suitable function for $f_Y(y)$ chosen, $f_{Y|X}(y|x)$ was approximated by fitting $f_Y(y)$ to particle landing data divided into slices over the x axis. Trends in the parameters of $f_Y(y)$ over x would be used to produce the function $f_{Y|X}(y|x)$.

With the mathematical form of each function $f_X(x)$ and $f_{Y|X}(y|x)$ having been determined, the final step is to produce a set of functions that determine the parameters of $f_X(x)$ and $f_{Y|X}(y|x)$ in terms of the scenario conditions. This was done by investigating trends in each parameter across sets of scenarios grouped by physical conditions to identify which physical conditions had the greatest impact, and how the physical conditions interact with each other to influence the parameters of these functions. Generalizations of these trends were tested by producing distributions using preliminary forms of the statistical model and comparing the shape of these distributions against FDS simulation results.

8.2.4 Model Performance

The performance of the final statistical model of firebrand landing was evaluated using similar statistical tools to those used in the evaluation of $f_X(x)$ and $f_{Y|X}(y|x)$. The overall performance of the model was very good, with small total discrepancies between statistical model predictions of particle landing density and FDS simulation results. The statistical model performed best for cases with larger fires, higher wind speeds, and higher HRRs. The model

performs similarly well in both forested and non-forested domains. Prediction quality was poorest for cases with small W:D ratios and low wind speeds, suggesting that the functions used to approximate the distribution parameters could be improved with more data on particle distribution in smaller and narrower fires. The computational resources required to run the statistical model are minimal, with a desktop PC able to produce a distribution of hundreds of thousands of points in just a few minutes.

8.3 Recommendations for Future Research

While the statistical model that has been developed over the course of this project performs well, the model could be expanded upon considerably. The high computational cost of CFD simulation limited the range of scenarios that could be feasibly completed and analyzed during this research project. Future model development efforts should incorporate simulation results for a much broader range of flame dimensions, HRRs, wind speeds, and domain types to produce a model with a broader range of application to real fire scenarios. In particular, the effect of different canopy heights, vegetation density profiles, and terrain topography should be investigated – such as the effect of domains with slopes, hills, and valleys on particle distribution.

In addition to this, data on firebrand geometries and masses from a broader range of vegetation types should be incorporated into the model. As the model is currently constructed, it would be simple to adjust the relative frequency of particles with different masses and geometries to produce a distribution for a different type of vegetation. However, if a vegetation species produced particles that are not well-approximated as cylindrical, spherical, or rectangular geometries, a new set of model fit parameters would need to be generated to model the distribution of these particles. With a broad set of data on particle mass, geometry, and generation rate characteristics for different vegetation types, the applicability of the model could be expanded considerably. Particular attention should be given to particles with especially high or low densities relative to the set of particles used to develop the current model.

Verification of the model against data obtained from field-scale controlled burning experiments that record firebrand landing positions would also be highly useful. Although it is difficult to obtain this data in practice, some experimental data on firebrand landing position from Project Vesta has been published, including plots of firebrand landing density that show similar characteristics to distributions obtained from FDS simulation in this research. An example of

these distributions are shown in Chapter 2: Figure 2-2 shows a clear example of the bimodal distribution seen in our simulations and model.

In terms of statistical methods, alternative techniques for modelling the distribution of particles could be tested. Modern machine learning models and AI techniques may be able to extract more subtle trends in particle distribution from the available data. These techniques may be used to either augment and improve the existing model, or to generate new models entirely. For the purpose of augmenting the existing model, machine learning algorithms may be used to find more optimal relationships between the scenario conditions (wind speed, HRR, W:D ratio, forest or open terrain) and the parameters of the functions $f_X(x)$ and $f_{Y|X}(y|x)$. Alternatively, an AI model could be trained on the dataset obtained from FDS simulations and tested against the existing statistical model. Research on developing Explainable AI models (XAI) [99] may allow for the results of an AI model to be understood from a physics perspective, producing a model with good performance characteristics that can also provide insight into the physics of particle transport.

Extending the model with a function that describes the likelihood of igniting a spot fire for given vegetation conditions would be a natural progression of this research. As the model currently uses a different set of parameters for $f_X(x)$ for each particle type based on particle mass and geometry, the model could be modified to predict the proportion of each type of particle likely to land in a given area. If each particle type can be assigned a function that describes the likelihood of igniting a fuel bed at a given distance (given the particles lose mass to combustion or cool over distance) it may be possible to adapt the model to produce a map of spot fire ignition probability over the domain.

References:

1. Bento-Gonçalves, A. and A. Vieira, *Wildfires in the wildland-urban interface: Key concepts and evaluation methodologies*. Science of the total environment, 2020. **707**: p. 135592.
2. Mell, W.E., et al., *The wildlandurban interface fire problem current approaches and research needs*. International Journal of Wildland Fire, 2010. **19**(2): p. 238-251.
3. Jolly, W.M., et al., *Climate-induced variations in global wildfire danger from 1979 to 2013*. Nature communications, 2015. **6**: p. 7537.
4. Huang, Y., S. Wu, and J.O. Kaplan, *Sensitivity of global wildfire occurrences to various factors in the context of global change*. Atmospheric Environment, 2015. **121**: p. 86-92.
5. IMAGE-team, *The IMAGE 2.2 Implementation of the SRES Scenarios. A Comprehensive Analysis of Emissions, Climate Change and Impacts in the 21st Century; CD-ROM 500110001 (former 481508018)*. 2006, National Institute for Public Health and the Environment Bilthoven, The
6. Boegelsack, N., et al., *A Critical Examination of the Relationship between Wildfires and Climate Change with Consideration of the Human Impact*. Journal of Environmental Protection,, 2018. **9**: p. 461-467.
7. Abatzoglou, J.T. and A.P. Williams, *Impact of anthropogenic climate change on wildfire across western US forests*. Proceedings of the National Academy of Sciences, 2016. **113**(42): p. 11770-11775.
8. Di Virgilio, G., et al., *Climate Change Increases the Potential for Extreme Wildfires*. Geophysical Research Letters, 2019. **46**(14): p. 8517-8526.
9. Mark Binskin, D.A.B., Professor Andrew Macintosh *Royal Commission into National Natural Disaster Arrangements Report*. 2020, Royal Commission.
10. Protection, C.D.o.F.a.F. *2020 Incident Archive*. 2020; Available from: <https://www.fire.ca.gov/incidents/2020/>.
11. Bureau of Meteorology, A. *Twelve-monthly mean temperature anomaly for Australia*. 2021; Available from: <http://www.bom.gov.au/jsp/awap/temp/archive.jsp?colour=colour&time=latest&step=0&map=meananom&month=12&period=12month&area=nat&year=2019>.
12. Bureau of Meteorology, A. *Annual climate statement 2023*. 2023; Available from: <http://www.bom.gov.au/climate/current/annual/aus/#tabs=Temperature>.
13. *California Heatwave Fits a Trend*. [cited 2023; Available from: <https://earthobservatory.nasa.gov/images/147256/california-heatwave-fits-a-trend>.
14. Sullivan, A.L., *Wildland surface fire spread modelling, 1990–2007. 2: Empirical and quasi-empirical models*. International Journal of Wildland Fire, 2009. **18**(4): p. 369-386.
15. McArthur, A.G., *Fire behaviour in eucalypt forests*, in *Pap. 9th Commonw. For. Conf., New Delhi*. 1967.
16. Rothermel, R.C., *A mathematical model for predicting fire spread in wildland fuels*. Res. Pap. INT-115. Ogden, UT: US Department of Agriculture, Intermountain Forest and Range Experiment Station. 40 p., 1972. **115**.
17. Cruz, M.G. and M.E. Alexander, *Uncertainty associated with model predictions of surface and crown fire rates of spread*. Environmental Modelling & Software, 2013. **47**: p. 16-28.
18. Ghaderi, M., M. Ghodrat, and J.J. Sharples, *LES Simulation of Wind-Driven Wildfire Interaction with Idealized Structures in the Wildland-Urban Interface*. Atmosphere, 2021. **12**(1): p. 21.

19. McGrattan, K., et al., *Fire Dynamics Simulator User's Guide*. 2021: NIST.
20. McGrattan, K., et al., *Fire Dynamics Simulator Technical Reference Guide Volume 3: Validation*. 2024, National Institute of Standards and Technology.
21. Wadhwani, R., et al., *Verification of a Lagrangian particle model for short-range firebrand transport*. Fire Safety Journal, 2017. **91**: p. 776-783.
22. Wadhwani, R., et al., *Firebrand transport from a novel firebrand generator: numerical simulation of laboratory experiments*. International Journal of Wildland Fire, 2022. **31**(6): p. 634-648.
23. Blanchi, R. and J. Leonard, *Investigation of Bushfire Attack Mechanisms Resulting in House Loss in the ACT Bushfire 2003, A Bushfire Cooperative Research Centre (CRC) Report*. 2005, CMIT Technical Report-2005.
24. Sydney, S.A., *AS 3959-2009: Construction of buildings in bushfire prone areas, Standards Australia: Sydney*. 2009.
25. Cruz, M.G., et al., *Guide to Rate of Fire Spread Models for Australian Vegetation*. 2015: CSIRO Land and Water Flagship, Canberra, ACT and AFAC, Melbourne, VIC.
26. Wang, H.-H., *Analysis on Downwind Distribution of Firebrands Sourced from a Wildland Fire*. Fire Technology, 2011. **47**(2): p. 321-340.
27. Luke, R.H. and A.G. McArthur, *Bushfires in Australia*. 1978: Australian Government Publishing Service for CSIRO.
28. Wadhwani, R., et al., *A review of firebrand studies on generation and transport*. Fire Safety Journal, 2022. **134**: p. 103674.
29. Ganteaume, A., et al., *Laboratory characterization of firebrands involved in spot fires*. Annals of Forest Science, 2011. **68**(3): p. 531-541.
30. Manzello, S.L., A. Maranghides, and W.E. Mell, *Firebrand generation from burning vegetation*. International Journal of Wildland Fire, 2007. **16**(4): p. 458-462.
31. Wickramasinghe, A., N. Khan, and K. Moinuddin, *Determining Firebrand Generation Rate Using Physics-Based Modelling from Experimental Studies through Inverse Analysis*. Fire, 2022. **5**(1): p. 6.
32. Bahrani, B., *Characterization of Firebrands Generated from Selected Vegetative Fuels in Wildland Fires*. 2020, The University of North Carolina at Charlotte.
33. Hudson, T.R., et al., *Effects of fuel morphology on ember generation characteristics at the tree scale*. International Journal of Wildland Fire, 2020. **29**(11): p. 1042-1051.
34. Ellis, P.F., *The effect of the aerodynamic behaviour of flakes of jarrah and karri bark on their potential as firebrands*. Journal of the Royal Society of Western Australia, 2010. **93**(1): p. 21-27.
35. Almeida, M., L. Porto, and D. Viegas, *Characterization of Firebrands Released From Different Burning Tree Species*. Frontiers in Mechanical Engineering, 2021. **7**.
36. Gould, J.S., et al., *Project Vesta: fire in dry eucalypt forest: fuel structure, fuel dynamics and fire behaviour*. 2008: Csiro Publishing.
37. Dean, W., *LXXII. The stream-line motion of fluid in a curved pipe (Second paper)*. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 1928. **5**(30): p. 673-695.
38. Thomas, J.C., et al., *Investigation of firebrand generation from an experimental fire: Development of a reliable data collection methodology*. Fire Safety Journal, 2017. **91**: p. 864-871.
39. Wadhwani, R., *Physics-based simulation of short-range spotting in wildfires*, in *College of Engineering and Sciences*. 2019, Victoria University: Melbourne.
40. Haider, A. and O. Levenspiel, *Drag coefficient and terminal velocity of spherical and nonspherical particles*. Powder technology, 1989. **58**(1): p. 63-70.



41. Tohidi, A. and N.B. Kaye, *Comprehensive wind tunnel experiments of lofting and downwind transport of non-combusting rod-like model firebrands during firebrand shower scenarios*. Fire Safety Journal, 2017. **90**: p. 95-111.
42. Tohidi, A., *Experimental and Numerical Modeling of Wildfire Spread via Fire Spotting*. 2016, Clemson University.
43. Hage KD, "On the dispersion of large particles from a 15-m source in the atmosphere". J Meteor 1961. **18**:534–539.
44. Tarifa, C.S., *Transport and combustion of firebrands*, U.S.D.O.A.F. SERVICE, Editor. 1967, Instituto Nacional de Tecnica Aeroespacial, Esteban Terradas: Madrid.
45. Wang, H. *Ember attack: its role in the destruction of houses during ACT bushfire in 2003*. in *Life in a fire-prone environment: translating science into practice. Bushfire Conf*. 2006.
46. Himoto, K. and T. Tanaka, *Transport Of Disk-shaped Firebrands In A Turbulent Boundary Layer*. Fire Safety Science, 2005. **8**: p. 433-444.
47. Sardoy, N., et al., *Numerical study of ground-level distribution of firebrands generated by line fires*. Combustion and Flame, 2008. **154**(3): p. 478-488.
48. Sardoy, N., et al., *Modeling transport and combustion of firebrands from burning trees*. Combustion and Flame, 2007. **150**(3): p. 151-169.
49. Noble, I., A. Gill, and G. Bary, *McArthur's fire-danger meters expressed as equations*. Australian Journal of Ecology, 1980. **5**(2): p. 201-203.
50. Burrows, N.D., *Experimental development of a fire management model for Jarrah (Eucalyptus Marginata Donn ex Sm.) forest*. 1994, Australian National University, Canberra.
51. Burrows, N., *Fire behaviour in jarrah forest fuels: 2. Field experiments*. 1999.
52. Tarifa, C.S., P.P. del Notario, and F.G. Moreno. *On the flight paths and lifetimes of burning particles of wood*. in *Symposium (International) on Combustion*. 1965. Elsevier.
53. Ellis, P.F., *The aerodynamic and combustion characteristics of eucalypt bark: a firebrand study*. 2000, Australian National University.
54. Raupach, M., *Similarity analysis of the interaction of bushfire plumes with ambient winds*. Mathematical and Computer Modelling, 1990. **13**(12): p. 113-121.
55. Hu, H.H., *Chapter 10 - Computational Fluid Dynamics*, in *Fluid Mechanics (Fifth Edition)*, P.K. Kundu, I.M. Cohen, and D.R. Dowling, Editors. 2012, Academic Press: Boston. p. 421-472.
56. Piomelli, U., *Large eddy simulations in 2030 and beyond*. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 2014. **372**(2022): p. 20130320.
57. Rehm, R.G. and H.R. Baum, *The Equations of Motion for Thermally Driven, Buoyant Flows*. J Res Natl Bur Stand (1977), 1978. **83**(3): p. 297-308.
58. Deardorff, J.W., *Stratocumulus-capped mixed layers derived from a three-dimensional model*. Boundary-layer meteorology, 1980. **18**: p. 495-527.
59. Pope, S.B., *Turbulent Flows*. 2000, Cambridge: Cambridge University Press.
60. McGrattan, K., et al., *Fire Dynamics Simulator Technical Reference Guide Volume 1: Mathematical Model*. 2015, National Institute of Standards and Technology: Gaithersburg, Maryland, USA. p. 145.
61. Courant, R., K. Friedrichs, and H. Lewy, *On the Partial Difference Equations of Mathematical Physics*. IBM Journal of Research and Development, 1967. **11**(2): p. 215-234.



62. Cunningham, P.G., S. L.; Hussaini, M. Y.; Linn, R. R., *Coherent vortical structures in numerical simulations of buoyant plumes from wildland fires*. International Journal of Wildland Fire, 2005.
63. McArthur, A.G., *Forest fire danger meter*. 1973: Forest Research Institute, Forestry & Timber Bureau.
64. Wadhvani, R., et al., *Kinetics of pyrolysis of litter materials from pine and eucalyptus forests*. Journal of Thermal Analysis and Calorimetry, 2017. **130**(3): p. 2035-2046.
65. Matthews, S.G., James & Mccaw, Lachie., *Simple models for predicting dead fuel moisture in eucalyptus forests*. International Journal of Wildland Fire., 2010. **19**. : p. 459-467.
66. Manzello, S.L., et al., *Mass and size distribution of firebrands generated from burning Korean pine (Pinus koraiensis) trees*. Fire and Materials: An International Journal, 2009. **33**(1): p. 21-31.
67. El Houssami, M., et al., *Experimental procedures characterising firebrand generation in wildland fires*. Fire technology, 2016. **52**(3): p. 731-751.
68. Jarrin, N., et al., *A synthetic-eddy-method for generating inflow conditions for large-eddy simulations*. International Journal of Heat and Fluid Flow, 2006. **27**(4): p. 585-593.
69. McGrattan, K., et al., *Fire dynamics simulator (Sixth Edition) user's guide*, in *Special publication 1019*. 2015, National Institute of Standards and Technology: Gaithersburg, Maryland, USA. p. 280.
70. Basu, S., Lacser, A., *A Cautionary Note on the Use of Monin–Obukhov Similarity Theory in Very High-Resolution Large-Eddy Simulations*. Boundary-Layer Meteorology, 2017.
71. Celik, I.B., U. Ghia, and P.J. Roache, *Procedure for estimation and reporting of uncertainty due to discretization in CFD applications*. Journal of Fluids Engineering, 2008. **130**(7).
72. Roache, P.J., *Perspective: A Method for Uniform Reporting of Grid Refinement Studies*. Journal of Fluids Engineering, 1994.
73. Forney, G.P., *Smokeview, A Tool for Visualizing Fire Dynamics Simulation Data, Volume I: User's Guide*. Vol. 1. 2024: National Institute of Standards and Technology.
74. Tory, K., *Models of Buoyant Plume Rise*. 2018, Bushfire and Natural Hazards CRC.
75. Barletta, A., *The Boussinesq approximation for buoyant flows*. Mechanics Research Communications, 2022. **124**: p. 103939.
76. Stull, R., *Meteorology for Scientists & Engineers, 3rd Edition*. 2011: Univ. of British Columbia.
77. Briggs, G.A., *Some recent analyses of plume rise observations*. 1970.
78. Fisher, B.E.A., et al., *Modelling plume rise and dispersion from pool fires*. Atmospheric Environment, 2001. **35**(12): p. 2101-2110.
79. Zonato, C., et al., *Plume rise of smoke coming from free burning fires*. Journal of Hazardous Materials, 1993. **34**(1): p. 69-79.
80. Haines, D.A. and M.C. Smith, *Three Types of Horizontal Vortices Observed in Wildland Mass and Crown Fires*. Journal of Applied Meteorology and Climatology, 1987. **26**(12): p. 1624-1637.
81. Forthofer, J.M. and S.L. Goodrick, *Review of vortices in wildland fire*. Journal of Combustion, 2011. **2011**(1): p. 984363.
82. Church, C.R., J.T. Snow, and J. Dessens, *Intense Atmospheric Vortices Associated with a 1000 MW Fire*. Bulletin of the American Meteorological Society, 1980. **61**(7): p. 682-694.



83. Shinohara, M. and S. Matsushima, *Formation of fire whirls: Experimental verification that a counter-rotating vortex pair is a possible origin of fire whirls*. Fire Safety Journal, 2012. **54**: p. 144-153.
84. Haines, D.A. and M.C. Smith, *Simulation of the Collapse of Bent-Over Vortex Pairs Observed in Wildland Fires*. Forest Science, 1992. **38**(1): p. 68-79.
85. Li, B., et al., *Numerical investigation of the flow characteristics around two tandem propane fires in a windy environment*. Fuel, 2021. **286**: p. 119344.
86. Lareau, N.P., et al., *Fire-Generated Tornadic Vortices*. Bulletin of the American Meteorological Society, 2022. **103**(5): p. E1296-E1320.
87. Barata, B.A.C., F.C. Martins, and J.C.F. Pereira, *LES of fire plumes subjected to crosswind inducing vertical vorticity*. Fire Safety Journal, 2024. **144**: p. 104112.
88. Christopher R. Church, J.T.S., and Jean Dessens, *Intense Atmospheric Vortices Associated with a 1000 MW Fire*. Bulletin of the American Meteorological Society. **67**(7): p. 682-694.
89. Nair, V., et al., *Counter rotating vortex pair structure in a reacting jet in crossflow*. Proceedings of the Combustion Institute, 2019. **37**(2): p. 1489-1496.
90. Chen, Y.-C., *A tutorial on kernel density estimation and recent advances*. Biostatistics & Epidemiology, 2017. **1**(1): p. 161-187.
91. Filkov, A., et al., *Investigation of firebrand production during prescribed fires conducted in a pine forest*. Proceedings of the Combustion Institute, 2017. **36**(2): p. 3263-3270.
92. Myung, I.J., *Tutorial on maximum likelihood estimation*. Journal of Mathematical Psychology, 2003. **47** (2003): p. 90-100.
93. Freund, J.E., *Mathematical Statistics, 5th Edition*. 1992: Prentice Hall International.
94. Akaike, H., *A new look at the statistical model identification*. IEEE Transactions on Automatic Control, 1974. **19**(6): p. 716-723.
95. Anderson, T.W., *On the Distribution of the Two-Sample Cramer-von Mises Criterion*. The Annals of Mathematical Statistics, 1962. **33**(3): p. 1148-1159.
96. Durante, F., J. Fernández-Sánchez, and C. Sempi, *A topological proof of Sklar's theorem*. Applied Mathematics Letters, 2013. **26**(9): p. 945-948.
97. Park, K.I., *Fundamentals of Probability and Stochastic Processes with Applications to Communications*. 2018: Springer Cham.
98. Hodson, T.O., *Root-mean-square error (RMSE) or mean absolute error (MAE): when to use them or not*. Geosci. Model Dev., 2022. **15**(14): p. 5481-5487.
99. Dwivedi, R., et al., *Explainable AI (XAI): Core Ideas, Techniques, and Solutions*. ACM Comput. Surv., 2023. **55**(9): p. Article 194.



Chapter 9 - Appendix I



This appendix contains overlaid kernel density estimate (KDE) plots of the distribution of particles obtained from FDS simulation and the distribution of particles produced by the statistical model. KDE plots are also presented for the distribution in the x and y axes separately. Summary statistics describing the level of agreement between the simulation and model results are presented in Table 4 of the Model Description and Performance chapter.























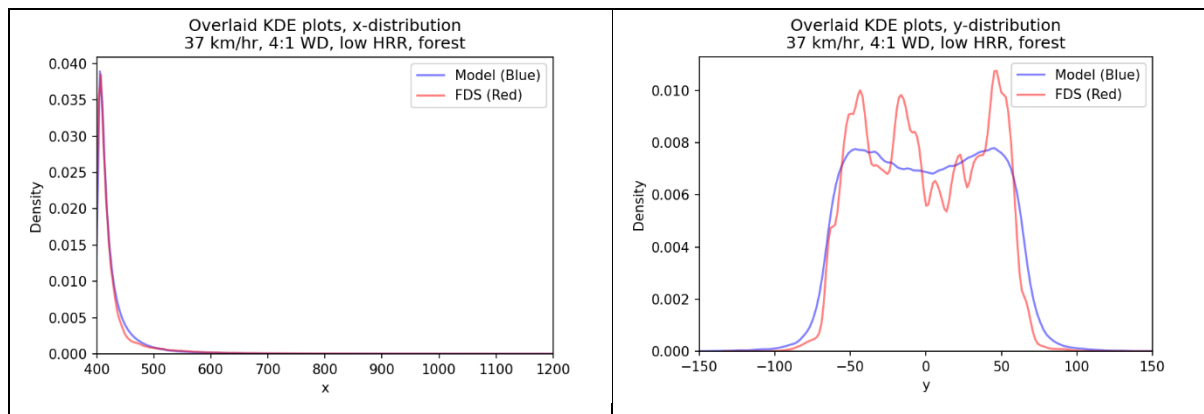





Overlaid KDE contours
37 km/hr, 2:1 WD, low HRR, open

































Overlaid KDE contours
46 km/hr, 1:1 WD, low HRR, open






















































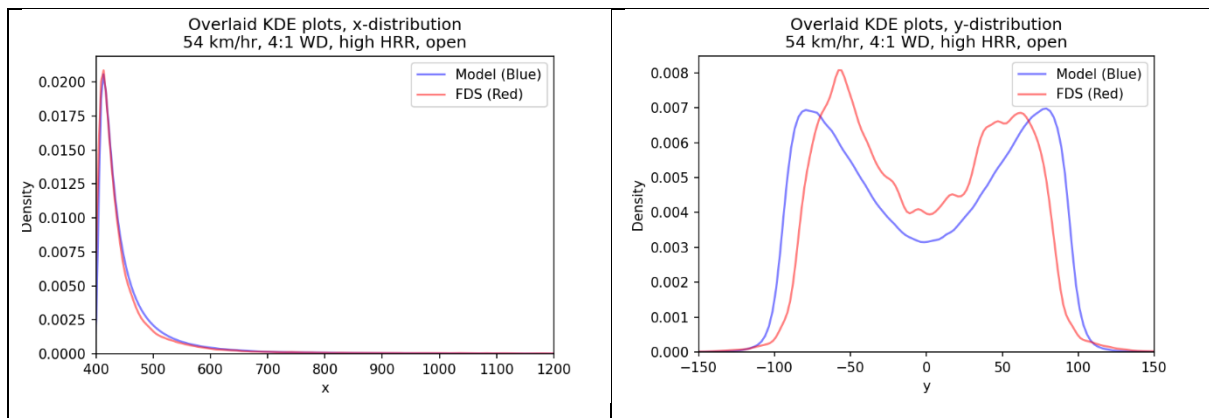

















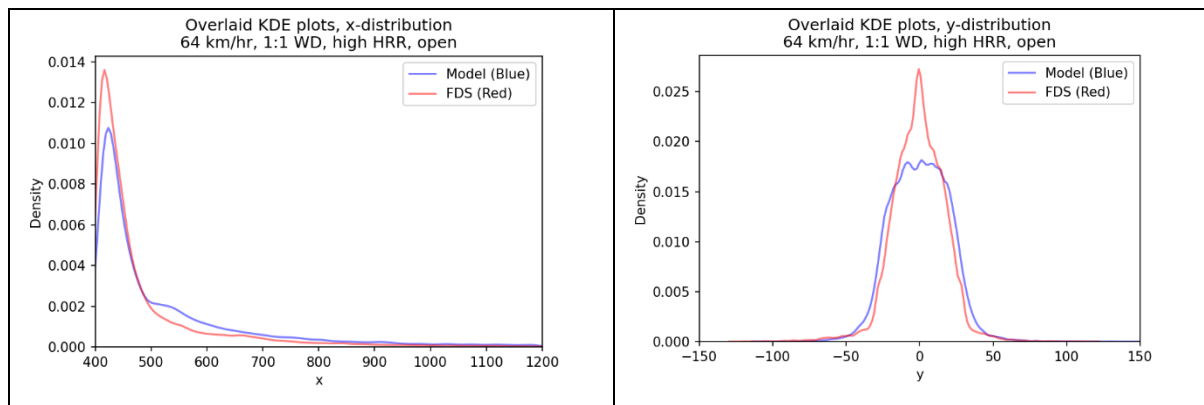









Overlaid KDE contours
54 km/hr, 4:1 WD, low HRR, open





























Overlaid KDE contours
64 km/hr, 2:1 WD, low HRR, open











Chapter 10 - Appendix II



This appendix contains graphs of the shape of the convective plume overlaid with a predicted plume shape based on the Zonato plume rise model. The plots are grouped into two categories: forested and open field cases.



Open Field Cases



These plots are shown in order of increasing wind speed and W:D ratio. Each plot has been assigned a letter (A to V) so that it can be easily referenced in the text.











































Forested Cases



Similar to the Open Field cases, these plots are shown in order of increasing wind speed and W:D ratio. Each plot has been assigned a letter so that it can be easily referenced in the text. One case (64km/hr, 2:1 W:D, Low HRR, Forested) had a corrupted data file that made it impossible to produce a plume rise graph, and has been omitted from this appendix.












































Chapter 11 - Appendix III

This appendix contains plots of the trajectory of a random sample of particles in each scenario. These plots are grouped into Open Field and Forested cases. Each plot is assigned a letter to identify it within its group, so that it can be referenced in the text easily.

Open Field Cases









G)



Particle trajectories, XY plane
46km/hr, 1:1 WD, high HRR, Open



H)


Particle trajectories, XY plane
46km/hr, 1:1 WD, low HRR, Open






K)

Particle trajectories, XY plane
46km/hr, 4:1 WD, high HRR, Open




L)

Particle trajectories, XY plane
46km/hr, 4:1 WD, low HRR, Open




M)

Particle trajectories, XY plane
54km/hr, 1:1 WD, high HRR, Open




N)

Particle trajectories, XY plane
54km/hr, 1:1 WD, low HRR, Open




O)

Particle trajectories, XY plane
54km/hr, 2:1 WD, high HRR, Open




P)

Particle trajectories, XY plane
54km/hr, 2:1 WD, low HRR, Open




Q)



Particle trajectories, XY plane
54km/hr, 4:1 WD, high HRR, Open



R)


Particle trajectories, XY plane
54km/hr, 4:1 WD, low HRR, Open






U)

Particle trajectories, XY plane
64km/hr, 2:1 WD, high HRR, Open







V)



Particle trajectories, XY plane
64km/hr, 2:1 WD, low HRR, Open







Forested Cases

















M)

Particle trajectories, XY plane
54km/hr, 1:1 WD, high HRR, Forest




N)

Particle trajectories, XY plane
54km/hr, 1:1 WD, low HRR, Forest




O)

Particle trajectories, XY plane
54km/hr, 2:1 WD, high HRR, Forest




P)

Particle trajectories, XY plane
54km/hr, 2:1 WD, low HRR, Forest




Q)



Particle trajectories, XY plane
54km/hr, 4:1 WD, high HRR, Forest





R)

Particle trajectories, XY plane
54km/hr, 4:1 WD, low HRR, Forest








W)

Particle trajectories, XY plane
64km/hr, 4:1 WD, high HRR, Forest



X)

Particle trajectories, XY plane
64km/hr, 4:1 WD, low HRR, Forest

