



THE UNIVERSITY  
*of* ADELAIDE

Developing a (Semi) Automatic Calibration Procedure for  
Cellular Automata based Land-use Models

Charles Peter Newland

BEng (Civil & Structural) Hons.

Thesis submitted to the University of Adelaide,  
Faculty of Engineering, Computer and Mathematical Sciences,  
School of Civil, Environmental and Mining Engineering,  
in fulfilment of the requirements for the degree of  
Doctor of Philosophy

Copyright © February 2018



## **Abstract**

Land-use change models are used to understand the wide-ranging impacts that land-use changes have on a region. Effective modelling of land-use changes must capture multiple, mutually influential drivers. A common framework for modelling land-use changes uses Cellular Automata (CA), which have seen a growth in application driven by the availability of generic modelling platforms, shifting the focus of research about Land-Use Cellular Automata (LUCA) models from development to application, with a particular focus on calibration.

Calibration of LUCA models is complex, as land-use change is a path-dependent process with uncertain outcomes captured by a number of model parameters. Of note are LUCA models that use a transition potential, which are traditionally calibrated using a manual approach, a process that is time-consuming and lacks objectivity. Hence, there has been a focus on the development of automatic calibration methods for these types of models. To automate calibration, metrics are used to capture two separate properties of performance: locational agreement, the match of pixels between simulated outputs and the corresponding observed data, and landscape pattern structure, the inferred realism of land-use change processes captured by the difference between the observed and simulated landscape patterns.

The primary objective of this research is to develop improved automatic calibration methods for transition potential based LUCA models. There are two common approaches, optimisation-based and process-specific. The major contributions of this body of work are the development of improved versions of each type of approach, and the development of a hybrid method combining the advantages of the two approaches.

First, a generic multi-objective optimisation framework for automatic calibration of transition potential LUCA models was developed in Paper 1 (Chapter 2) that allows for the exploration of trade-offs between the model performance objectives. Second, a process-specific semi-automatic calibration method that integrates objective analysis with discursive input to facilitate efficient calibration of neighbourhood rules (the main calibration parameter for this type of model) within a limited computational budget was developed in Paper 2 (Chapter 3). Finally, a generic framework for hybrid automatic calibration, which integrates domain knowledge into a multi-objective optimisation approach, was developed in Paper 3 (Chapter 4). The utility of each method was demonstrated via case study applications, showing promising potential for future applications of LUCA models to support long term planning and policy development.

## **Statement of originality**

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I acknowledge that copyright of published works contained within this thesis resides with the copyright holder(s) of those works.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Charles Peter Newland

(name)

28<sup>th</sup> Feb 2018

(signature)

(date)



## **Acknowledgements**

First, I would like to thank my PhD supervisors, Professor Holger Maier, Dr Aaron Zecchin, and Adjunct Associate Professor Hedwig van Delden, for their supervision, support, and tireless dedication throughout my PhD research. I would like to thank Professor Holger Maier for his guidance and enthusiasm, always driving my research forward. I would like to thank Dr Aaron Zecchin for his insight, and appreciation for detail. I would like to thank Adjunct Associate Professor Hedwig van Delden for her passion and commitment to the quality of my research. I would also like to thank Hedwig and her family, Roel, Elena and Laurens, for hosting me during the portions of my PhD candidature spent in Maastricht, The Netherlands.

I would like to acknowledge the tireless efforts of Research Associate Jeffrey Newman, who was instrumental in the software development stages of this research.

I am grateful for the primary funding I received via an Australian Postgraduate Award, that the supplementary funding provided by the Bushfire and Natural Hazards Cooperative Research Centre. I would like to acknowledge the Bushfire and Natural Hazards Cooperative Research Centre for providing opportunities for professional development, and to share my work with fellow researchers and practitioners.

I would like to thank my fellow research colleagues Sam Culley and Cameron McPhail for their support throughout my PhD, always helping me to see things more clearly.

To my partner, Rachel Evans, thank you for always being there for me with love, support, and understanding during my PhD. I am truly blessed to have you.

Finally, I wish to thank my family, my sisters Stephanie and Elizabeth, and my parents Peter and Rosemary Newland. We all went through a lot during my PhD, and it is a testament to the love, support, and unwavering spirit of our family that I was able to get through it all and complete my PhD. For that, I am truly thankful.

## Table of Contents

1	Introduction .....	1
1.1	Research objectives .....	4
1.2	Thesis overview.....	6
2	Multi-objective optimisation framework for calibration of Cellular Automata land-use models.....	9
2.1	Introduction .....	13
2.2	Proposed multi-objective optimisation-based calibration framework .....	15
2.3	Case study application of proposed framework .....	24
2.4	Results .....	34
2.5	Discussion .....	47
2.6	Conclusions and recommendations .....	50
2.7	Acknowledgements .....	51
2.8	Supplementary material.....	52
3	Empirically derived method and software for semi-automatic calibration of Cellular Automata land-use models.....	63
3.1	Introduction .....	67
3.2	Complexity of the calibration problem .....	69
3.3	Proposed approach .....	74
3.4	Application of proposed approach .....	87
3.5	Results and discussion.....	97
3.6	Summary and conclusions.....	108
3.7	Acknowledgements .....	109
3.8	Supplementary material.....	110
4	A hybrid (semi) automatic calibration method for Cellular Automata land-use models: Combining evolutionary algorithms with process understanding .....	139
4.1	Introduction .....	143

4.2	Proposed approach to integrating process understanding into optimisation-based automatic calibration.....	145
4.3	Case study .....	150
4.4	Results and discussion.....	159
4.5	Conclusions .....	182
4.6	Acknowledgements .....	183
4.7	Supplementary material.....	184
5	Conclusions .....	193
5.1	Research contribution.....	193
5.2	Limitations .....	196
5.3	Future work .....	197
6	References .....	199

## List of Figures

Figure 1.1. Research objectives .....	5
Figure 2.1. Proposed multi-objective optimization-based calibration framework for LUCA model with multiple dynamic land-use classes.....	16
Figure 2.2. Population-based metaheuristic for optimisation-based LUCA model parameter adjustment process .....	21
Figure 2.3. Location of the Randstad region and rasterized map for 1989.....	24
Figure 2.4. Parameterisation of a neighbourhood effect using exponential decay function, a locus point and critical distance .....	28
Figure 2.5. Optimisation process performed by NSGA-II, the parent population (P) is modified to generate a child population (Q), then recombined and modified further to produce the parent population (P') of the next generation .....	29
Figure 2.6. Composite Pareto front from the five optimisation runs, consisting of 77 solutions for parameter sets generating non-dominated objectives. The solutions are ordered by ascending FKS value (see colour bar) where the scatter of points arises from 10 independent model runs for each parameter set. The solid black line shows the averaged objective values over the 10 independent model runs. The dashed black lines show the calculated benchmark metric values for the growing clusters neutral model (lower) and random constraint match neutral model (higher). The five black dots show the solutions that are analysed in detail. ...	35
Figure 2.7. Comparison of selected simulated output maps corresponding to Pareto front solutions listed in Table 2.1 with data (a-1 and a-2). The maps are organised by increasing FKS values. ....	38
Figure 2.8. Variation of inertia parameter values for self-influence neighbourhood rules across Pareto front solutions .....	41
Figure 2.9. Variation of the conversion parameter influence values across Pareto front solutions for transitions from all land-use classes to the class residential .....	42
Figure 2.10. Variation of the neighbourhood rule tails across Pareto front solutions for the influence of the class residential in the neighbourhood of the class residential .....	43
Figure 2.11. Variation of the a parameter value across the Pareto front solutions, ordered by increasing FKS, for the influence of greenhouses in the neighbourhood of socio-cultural uses .....	45
Figure 2.12. Accessibility distance-decay parameter values across Pareto front solutions for motorways.....	46

Figure 2.13. Accessibility distance-decay parameter values across Pareto front solutions for other roads.....	46
Figure 2.14. Class level absolute clumpiness error variation across the Pareto front for five selected solutions .....	48
Figure. 3.1. Example of a transition potential allocation mechanism for a CA land-use model. Starting from an initial land-use map, the model allocates a transition potential to each cell (the gridded map) for each type of land-use based on a set of parametric maps, neighbourhood rules, and a stochastic component. The model then allocates a land-use class to each cell based on the potential and demand for each year (t). This process is repeated until the final year of the simulation (t=T), when the output map for that year is produced. ....	70
Figure 3.2. An example of how different neighbourhood rules and types of interactions are used to determine neighbourhood potential (in a single plane to the right of the cell marked with an x), an inertia point and self-influence tail, which influences the potential for the cell of interest to remain residential (red), and a conversion point and cross-influence tail, which influences the potential for the cell to transition to industrial (purple) .....	73
Figure 3.3. Conceptual outline of proposed methodology for semi-automatic process-specific neighbourhood rule calibration .....	75
Figure 3.4. Proposed interaction elimination, based on empirical evaluation and significance testing.....	76
Figure 3.5. An example tally of the frequency of cells of class j at distance d equal to 1 from a given land-use class transition (bottom row), as compared to the frequency of cells of class j at distance d equal to 1 from all cells (top row). For case a, the distributions appear similar, whereas for case b, the distributions are skewed to different extremes, suggesting differences that are statistically significant.....	80
Figure 3.6. Iterative procedure for performing coarse parameter adjustment where meta-parameters are tuned. ....	85
Figure 3.7. Iterative procedure for performing fine parameter adjustment via a line-search algorithm applied to individual neighbourhood weighting parameters .....	86
Figure 3.8. 1990 land-use maps of the four European case studies. ....	89
Figure 3.9. Number of interactions considered for conversion points and cross-influence tails before and after the interaction elimination stage for each case study .....	91
Figure 3.10. Example of the different behaviour observed in the objective space for FKS and AWCE when tuning a meta-parameter: convergent (a), where the selected point optimises both metrics, and trade-off (b), where improved performance in one objective results in reduced	

performance in the other. A trade-off requires user interpretation to preference a certain trade-off between objectives. ....94

Figure 3.11. Example of the different behaviour observed in the objective space for FK and AWCE during fine adjustment of the Lisbon case study, based on a preference for a certain objective. A preference tends to strongly bias the results to a certain metric, and a balanced preference can lead to improvement in both metrics. ....95

Figure 3.12. Average metric values for 50 model replicates compared to two sets of benchmark model metrics using statistical significance testing at the 95% confidence level for the calibration period. Cells are coloured by performance compared to the benchmark, green indicates the average is significantly superior to both benchmarks, yellow indicates the average is significantly superior to one benchmark, and red indicates the average is inferior to both benchmarks. ....98

Figure 3.13. Average metric values for 50 model replicates compared to two sets of benchmark model metrics using statistical significance testing at the 95% confidence level for the validation period. Cells are coloured by performance compared to the benchmark green indicates the average is significantly superior to both benchmarks, yellow indicates the average is significantly superior to one benchmark, and red indicates the average is inferior to both benchmarks. ....100

Figure 3.14. Berlin data compared with simulated output for the calibration period for best performing objective configuration. (a) is the data map, (b) is the simulated output map, (c) is the agreement of the class residential (red) between maps (a) and (b), and (d) is the agreement of the class industry & commerce (purple) between maps (a) and (b). ....102

Figure 3.15. Budapest data compared with simulated output for the calibration period for best performing objective configuration. (a) is the data map, (b) is the simulated output map, (c) is the agreement of the class residential (red) between maps (a) and (b), and (d) is the agreement of the class industry & commerce (purple) between maps (a) and (b). ....103

Figure 3.16. Lisbon data compared with simulated output for the calibration period for best performing objective configuration. (a) is the data map, (b) is the simulated output map, (c) is the agreement of the class residential (red) between maps (a) and (b), and (d) is the agreement of the class industry & commerce (purple) between maps (a) and (b). ....104

Figure 3.17. Madrid data compared with simulated output for the calibration period for best performing objective configuration. (a) is the data map, (b) is the simulated output map, (c) is the agreement of the class residential (red) between maps (a) and (b), and (d) is the agreement of the class industry & commerce (purple) between maps (a) and (b). ....105

Figure 3.18. Comparison of simulated output for different Madrid configurations that outperformed benchmarks. (a) is the data map, (b) is the simulated output for a balanced preference with FK and AWCE as objectives, and (c) is the simulated output for a balanced preference with FK, FKS and AWCE as objectives .....	106
Figure 3.19. Comparison of inertia point influence values across classes for the case studies .....	107
Figure 3.20. Comparison of self-influence tail parameter values across classes for each case study.....	108
Figure 4.1. Conceptual framework for proposed approach to integrating process understanding into optimisation-based automatic calibration .....	147
Figure 4.2. Comparison of triangular distributions with varying limits for sampling of parameter values .....	149
Figure 4.3. Location of Madrid region and rasterised land-use map for 1990 .....	151
Figure 4.4. Parameterisation of neighbourhood rule using an exponential decay function...	152
Figure 4.5. Objective performance of the output from the seeded method for the calibration period. Solutions analysed in more detail are highlighted in black. The degree of shading of the dots indicates the AWCE performance, with lighter dots corresponding to lower (better) AWCE values. The solid and dashed lines correspond to the growing clusters and random constraint match performance benchmark, respectively (benchmark error values for AWCE are not shown as they were significantly higher than the plot scale). .....	160
Figure 4.6. Objective performance of the output from the seeded method for the validation period. The degree of shading of the dots indicates the AWCE performance, with lighter dots corresponding to lower (better) AWCE values. The solid and dashed lines correspond to the growing clusters and random constraint match performance benchmark, respectively. ....	161
Figure 4.7. Simulated output maps for the calibration period for five selected solutions for the seeded method. (a-1) is the data map for 2000. The remaining maps are ordered into rows by decreasing FKS value, from (b) the output for Solution 1 (best FKS), to (f) the output for Solution 5 (best AWCE). The first column (e.g. (b-1)) shows the simulated output map for the selected solution, the second column (e.g. (b-2)) shows the agreement of the class residential between the data and corresponding simulated output, and the third column (e.g. (b-3)) shows the agreement of the class industry & commerce between the data and corresponding simulated output .....	164
Figure 4.8. Comparison of inertia point influence values per class for seeded solutions.....	167

Figure 4.9. Comparison of conversion point influence values per class for transitions to the class residential for the seeded solutions .....	168
Figure 4.10. Comparison of conversion point influence values per class for transitions to the class permanent crops for the seeded solutions. ....	169
Figure 4.11. Comparison of neighbourhood influence tails for urban classes to the class residential for the seeded solutions .....	170
Figure 4.12. Comparison of accessibility parameter values for classes residential, industry & commerce, and recreation areas for the seeded solutions .....	171
Figure 4.13. Comparison of objective performance for the calibration data for the optimisation methods run for 310 generations, and the process specific solutions. (a) a 3-D plot of the objective space; (b) a cross-section of the bi-objective space for FK vs AWCE; (c) a cross section of the bi-objective space for FKS vs AWCE; and (d) a cross section of the bi-objective space for FKS vs FK. Solid and dashed lines indicate the growing clusters and random constraint match neutral model benchmarks respectively. Black dots and diamonds show analysed solutions for the seeded and unseeded methods respectively. ....	174
Figure 4.14. Comparison of objective performance for the validation data for the optimisation methods run for 310 generations and the process specific method: (a) a 3d plot of the objective space; (b) a cross-section of the bi-objective space for FK vs AWCE; (c) a cross section of the bi-objective space for FKS vs AWCE; and (d) a cross section of the bi-objective space for FKS vs FK. Solid and dashed lines indicate the growing clusters and random constraint match neutral model benchmark respectively .....	176
Figure 4.15. Simulated output maps for the calibration period for five selected solutions for the unseeded method. (a-1) in the data map for 2000. The remaining maps are ordered into rows by decreasing FKS value. The first column (e.g. (b-1)) shows the simulated output map for the selected solution, the second column (e.g. (b-2)) shows the agreement of the class residential between the data and corresponding simulated output, and the third column (e.g. (b-3)) shows the agreement of the class industry & commerce between the data and corresponding simulated output .....	179
Figure 4.16. Simulated output maps for the process-specific calibration method compared with the calibration data (a-1), that resulted in better locational agreement performance (b-1) to (b-3) and better landscape pattern structure performance (c-1) to (c-3) for the calibration period. Maps in the middle column (e.g. (b-2)) show the agreement of the class residential between the simulated output and data. Maps in the right-most column (e.g. (b-3)) show the agreement of the class industry & commerce between the simulated output and data.....	180



## List of Tables

Table 2.1. Metric values of Pareto front solutions for which simulated output maps were evaluated .....	36
Table 2.2. Variation of conversion and influence at distance 1 of Pareto front solutions for conversion from greenhouses to socio-cultural uses and from socio-cultural uses to greenhouses .....	44
Table 2.3. A comparison of FKS, Kappa, and Fuzzy Kappa values across five selected Pareto front solutions .....	49
Table 3.1. Example contingency table, populated by logging the land-use class in each cell between two maps .....	78
Table 3.2. Summary of empirical measures used to categorise different interaction types for the proposed categorisation .....	83
Table 3.3. Thresholds used for inclusion in calibration .....	90
Table 3.4. Thresholds used for categorisation of parameters as high, medium or low for each interaction type .....	91
Table 3.5. Ranges and step sizes tested for meta-parameter sampling .....	92
Table 3.6. Combinations of metrics used as objectives for calibration .....	93
Table 3.7. Weighting used for each metric depending on objective preference .....	96
Table 4.1. Proportion of solution types used for initial population .....	156
Table 4.2. Benchmark metric values for different benchmark models .....	157
Table 4.3. Comparison of percentage of Pareto front solutions outperforming benchmark performance metrics and hyper-volume for seeded and unseeded calibration methods .....	182



# 1 Introduction

Land-use change models are used to understand the wide-ranging impacts that land-use change processes, such as urbanisation, deforestation and agricultural intensification, have on a region. As land-use changes are generally driven by policy and spatial planning, land-use change models are increasingly being used as part of decision support systems in the design and development of land-use change policy (Van Delden and Hurkens, 2011). Land-use change models have been applied to a diverse range of planning problems, including river basin management (Van Delden et al., 2007), developing sustainable agricultural practises (Connor et al., 2015, Van Delden et al., 2010, Murray-Rust et al., 2014c), exploring sustainable regional development plans (Rutledge et al., 2008) and urban growth planning (Berberoğlu et al., 2016, Chaudhuri and Clarke, 2013a).

Effectively modelling land-use changes requires capturing multiple, mutually influential bio-physical and socio-economic drivers to generate a realistic output (Lambin et al., 2001, Wang et al., 2011a). A common approach for modelling land-use changes uses Cellular Automata (CA), which aim to simulate the aggregate behaviour of multiple change agents, allowing for the exploration of large areas without the need for detailed data of actor behaviour (Hewitt et al., 2014). Their relative simplicity and intuitiveness have made CA a popular framework for modelling land-use changes (Santé et al., 2010).

Historically, CA were proposed to model geographic systems, such as land-use changes, by Tobler (1979) because they are an effective means of capturing the influence of spatial dynamics, a major driver of land-use change. Further investigation proposed that, with certain relaxations to the conventional CA structure (Couclelis, 1985), CA were capable of generating fractal patterns consistent with urban evolution (Batty and Longley, 1994, White and Engelen, 1993a). This led to the development of three major operational Land-Use Cellular Automata (LUCA) models (García et al., 2012), which vary based on how land-use changes are allocated: the SLEUTH pattern extrapolation model of Clarke et al. (1997) that allocates transitions by replicating different forms of urban growth and extrapolating each form based on the observed frequency, the development probability model of Wu (2002) that allocates land-use transitions based on a probability calculated from global and local factors, and the transition potential model of White and Engelen (1993c) that allocates land-use transitions based on the quantified potential ability of specific cells to support each particular land-use class.

The development of operational LUCA models led to an increase in the application of models that use the same structure. This growth has also been driven by the availability of generic modelling platforms SLEUTH and Metronamica (Van Delden and Hurkens, 2011), a transition potential model derived from White and Engelen (1993c). With consistent, well tested model architecture, and the availability of generic platforms allowing for direct application, research focus has shifted from LUCA model development to model application, with a particular focus on calibration.

Calibration of LUCA models is the process of determining a model parameter set, through the initial setting of model parameters, the iterative adjustment of these parameters based on comparison of the model output with observations, and the selection of a final parameter set, for application to a specific case study for long-term scenario analysis (Van Vliet et al., 2016). The iterative adjustment stage of LUCA model calibration is complex, as land-use change is a path-dependent process with uncertain outcomes (Brown et al., 2005) captured by a number of model parameters. The complexity of the calibration has led to a focus on automating the calibration procedure to make LUCA models more accessible and straightforward to apply for decision makers. Both the SLEUTH pattern extrapolation model (Silva and Clarke, 2002) and the development probability type model (Wu, 2002) have had (semi) automatic calibration methods since their inception, and both types of models have seen numerous developments in automatic calibration approaches (Clarke, 2018, Şalap-Ayça et al., 2018, Cao et al., 2014, Mustafa et al., 2018). By contrast, transition potential models are traditionally calibrated using a manual approach (White et al., 1997), which continues to be a common method of calibration (García et al., 2012, Van Delden et al., 2012), although several automatic calibration methods have been developed (Straatman et al., 2004, García et al., 2013, Van Vliet et al., 2013b, Blečić et al., 2015).

Manual calibration has remained common for transition potential LUCA models because it integrates discursive knowledge into the calibration process and ensures that results are consistent with process understanding. This is valuable because of the complexity of transition potential based LUCA models, which are commonly implemented with multiple dynamic land-use classes and use a wider neighbourhood for the consideration of spatial dynamics. Transition potential models extend the size of the neighbourhood of cells considered beyond a conventional CA model (the four or eight immediately adjacent cells) to all cells within a certain cellular radius, which requires the use of neighbourhood rules to characterise the influence different land-use classes exert on each other relative to proximity. These are the

main calibration parameters of this type of LUCA model (Engelen and White, 2008). Hence, transition potential LUCA models have a parameter space with much higher dimensionality (i.e. the number of parameters and the range of values they can take during the adjustment process) compared with other LUCA models, which makes calibration more complex. Despite the advantages of manual calibration, it is time-consuming, lacks objectivity, and is difficult to repeat (Jafarnejhad et al., 2016, García et al., 2013). Hence, there is an increasing focus on the development of automatic calibration methods for transition potential LUCA models (Van Vliet et al., 2013b).

Automating LUCA model calibration requires objective measures of LUCA model performance. Due to the path dependence and uncertainty associated with the process of land-use change, two separate properties of LUCA model performance should be considered: (i) locational agreement, alternatively termed cell-by-cell agreement (Hagen-Zanker, 2009), which is the match of pixels between simulated outputs and the corresponding observed data (Van Vliet et al., 2013b, Hagen-Zanker, 2009), and (ii) landscape pattern structure, which is the inferred realism of land-use change processes captured by the difference between the observed and simulated landscape patterns (Engelen and White, 2008). These have been considered to different extents in previous automatic calibration approaches, as outlined below.

There have been two approaches developed for the automatic calibration of transition potential based LUCA models. The first approach is optimisation-based, aiming to identify the best set of model parameters that optimise the measures of model performance used. Optimisation-based approaches are used because they are able to find (near) globally optimal solutions for highly complex (e.g. non-linear, non-convex) problems, and it is straightforward to link them with existing simulation models, such as LUCA models, without the need for problem simplification (Maier et al., 2015). The generic nature and advantages of optimisation has led to multiple applications to calibrate other forms of LUCA models (Feng et al., 2011, Liu et al., 2012, Clarke, 2018, Li et al., 2013).

As LUCA models have two distinct aspects of model performance, automatic calibration can be considered a multi-objective optimisation problem (Hagen-Zanker, 2008). However, at present, multi-objective optimisation has only been applied to automatically calibrate LUCA urban growth models implemented with a single (urban) active land-use class (Trunfio, 2006, Cao et al., 2014). In contrast, studies that use optimisation approaches for the calibration of LUCA models with multiple dynamic land-use classes have only considered a single objective.

For example, Blečić et al. (2015) considered only the locational agreement element of performance for automatic calibration, and while García et al. (2013) used both locational agreement and landscape pattern structure metrics, these were combined into a single objective during the optimisation process, meaning important trade-offs between locational agreement and landscape pattern structure could not be examined. Both studies also generated only one possible model parametrisation for future scenario analysis, limiting the ability to understand how calibrated parameters are potentially influenced by the metrics used for optimisation.

The second type of approach for automatic calibration of transition potential based LUCA models is process-specific, developed to generate a single calibrated set of neighbourhood rules using a single calibration objective as efficiently as possible. Such approaches generally attempt to mimic a manual calibration procedure to generate a set of neighbourhood rules that is consistent with process understanding. Despite a focus on neighbourhood rules, previous process-specific methods have not fully utilised discursive knowledge (White et al., 1997, Hagoort et al., 2008) to generate neighbourhood rules that are consistent with process understanding (Van Vliet et al., 2013b), and do not necessarily focus on the most important neighbourhood rules during the automatic calibration process (Straatman et al., 2004, Maas et al., 2005). Previous methods have also only used a single metric of performance, not considering the implementation of multiple metrics to capture the two aspects of LUCA model performance, and how these competing objectives impact on the resulting model.

## **1.1 Research objectives**

The objectives of this body of research are summarised in Figure 1.1. As shown, this research is divided between a set of primary objectives and secondary objectives. The overall aim of this research is to develop improved (semi) automatic calibration methods for transition potential based LUCA models. This development will ultimately aid in the evaluation of long-term policy and spatial planning by making the application of transition potential based LUCA models more straightforward, objective, and efficient. Given that there are two distinct conceptual approaches to automatic calibration, optimisation-based and process-specific, this study has three primary objectives that fall under this overall aim, as shown.

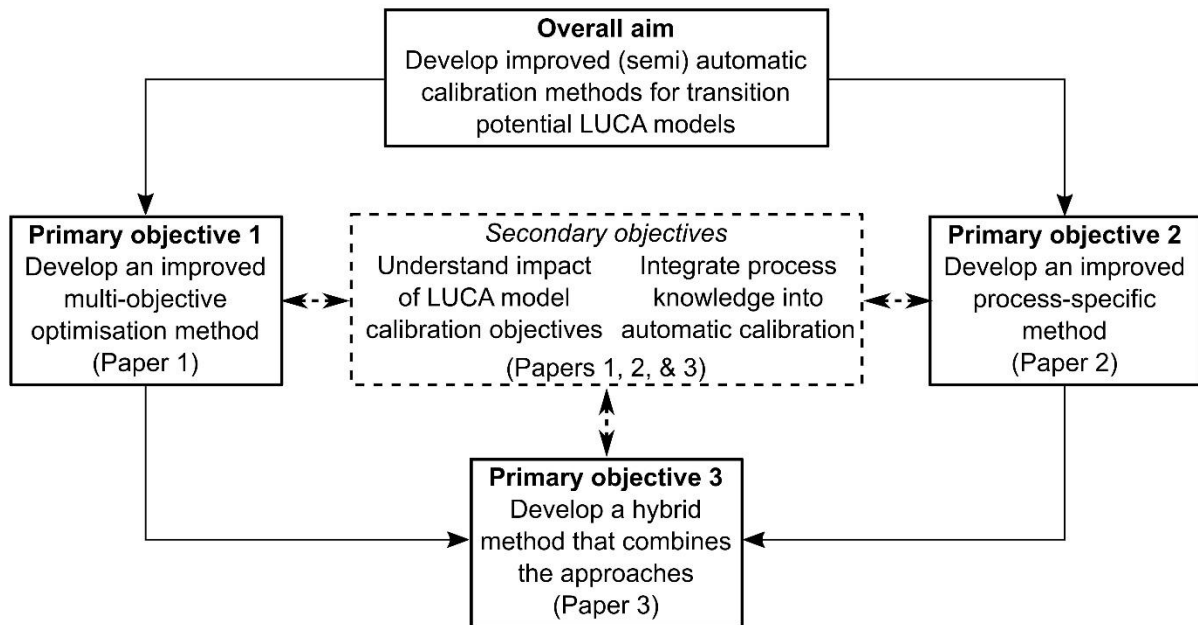


Figure 1.1. Research objectives

**Primary objective 1:** To develop an improved multi-objective optimisation framework that automates the calibration of LUCA models with multiple dynamic land-use classes, enabling the identification of multiple model parameterisations that could be suitable for long-term scenario analysis, and to demonstrate the application of the framework to a real-world case-study (Paper 1);

**Primary objective 2:** To develop an improved calibration method that utilises process knowledge about meaningful interactions to facilitate efficient automatic calibration that allows for consideration of both aspects of LUCA model performance for a limited computational budget, to demonstrate the utility of such an approach, and to determine the impact of a preference for different objectives on the resultant calibrated model (Paper 2);

**Primary objective 3:** To introduce a hybrid method for integrating the results of a process-specific (semi) automatic calibration method into an optimisation-based approach, and to investigate the utility of this approach compared with the standard formulation of the other two methods (Paper 3).

In addition to the primary objectives outlined, this research has secondary objectives that are naturally aligned with the development of automatic calibration methods, and will further enhance automatic calibration methods:

*Secondary objective 1:* To enhance understanding of how the metrics used to quantify the separate properties of LUCA model performance, locational agreement and landscape pattern structure, impact the resulting model parameterisation and simulated output obtained at the conclusion of automatic calibration (Papers 1, 2, and 3).

*Secondary objective 2:* To integrate process understanding into the application and analysis of automatic calibration methods to gain insights into how this knowledge can improve automatic calibration, and gain a deeper understanding of LUCA model behaviour (Papers 1, 2, and 3).

## **1.2 Thesis overview**

The remainder of this thesis is organised with the main contributions presented in Chapters 2 to 4. Each chapter is presented as a technical journal paper, with Chapter 2 a published paper in *Environmental Modelling and Software* and Chapter 3 a revised version of a paper submitted to *Environmental Modelling and Software* in response to reviewer comments.

**Chapter 2** presents a generic multi-objective optimisation framework for the automatic calibration of a LUCA model with multiple dynamic land-use classes. The generic nature of the framework facilitates substitution of the different components (e.g. land-use model, metrics) and a detailed exploration of the trade-off between the LUCA model performance objectives. The utility of the proposed framework is demonstrated via application to a case study of the Randstad region in the Netherlands.

**Chapter 3** presents an original process-specific calibration method for the (semi) automatic calibration of LUCA models. The proposed method is divided into two components, first simplifying the problem to the core set of neighbourhood rule parameters before calibrating these parameters using a two-stage procedure to ensure calibration is as efficient as possible. The utility of the proposed method is demonstrated via application to four case studies of major European cities.

**Chapter 4** presents a method developed to combine the two approaches for (semi) automatic calibration of LUCA models that utilises the benefits of both. The method allows for the generic multi-objective optimisation framework developed in Chapter 2 to use the output of the process-specific calibration approach developed in Chapter 3 so that the optimisation process starts in promising regions of the parameter space for efficient automatic calibration that allows for exploration of the trade-off between the model performance objectives used. The utility of the proposed approach is demonstrated via application to a case-study of Madrid, Spain.



A synthesis of the research presented in this thesis is presented in **Chapter 5**, which summarises the main research contributions of this body of work, the research limitations, and future directions for research in this field.



## **2 Multi-objective optimisation framework for calibration of Cellular Automata land-use models**

Newland, C.P., Maier, H.R., Zecchin, A.C., Newman, J.P. & Van Delden, H. 2018. Multi-objective optimisation framework for calibration of Cellular Automata land-use models. *Environmental Modelling & Software*, 100, 75-200. (Adapted for consistency).

## Statement of Authorship

Title of Paper	Multi-objective optimisation framework for calibration of Cellular Automata land-use models
Publication Status	<input checked="" type="checkbox"/> Published <input type="checkbox"/> Accepted for publication <input type="checkbox"/> Submitted for publication <input type="checkbox"/> Unpublished and Un-submitted work written in manuscript style
Publication Details	Newland, C.P. Maier, H.R., Zecchin, A.C., Newman, J.P. & Van Delden, H. 2018. Multi-objective optimisation framework for calibration of Cellular Automata land-use models. <i>Environmental Modelling &amp; Software</i> , 100, 75-200.

## Principal Author

Name of Principal Author (Candidate)	Charles P. Newland		
Contribution to the Paper	Designed scope of study and experimental procedure, developed software, conducted experiments, performed analysis of results, wrote manuscript and acted as corresponding author.		
Overall percentage (%)	70%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.		
Signature		Date	28 <sup>th</sup> Feb 2018

## Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

- i. The candidate's stated contribution to the publication is accurate (as detailed above);
- ii. Permission is granted for the candidate to include the publication in the thesis; and
- iii. The sum of all co-author's contributions is equal to 100% less the candidate's stated contribution.

Name of Co-author	Holger R. Maier		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript and response to reviewers.		
Signature		Date	28 <sup>th</sup> Feb 2018

Name of Co-author	Aaron C. Zecchin		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript and response to reviewers.		
Signature		Date	28 <sup>th</sup> Feb 2018

Name of Co-author	Jeffrey P. Newman		
Contribution to the Paper	Assisted with software development for super-computer and conducting experiments. Reviewed manuscript.		
Signature		Date	28 <sup>th</sup> Feb 2018

Name of Co-author	Hedwig Van Delden		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript and response to reviewers.		
Signature		Date	28 <sup>th</sup> Feb 2018

# **Multi-objective optimisation framework for calibration of Cellular Automata land-use models**

## **Abstract**

Modelling of land-use change plays an important role in many areas of environmental planning. However, land-use change models remain challenging to calibrate, as they contain many sensitive parameters, making the calibration process time-consuming. We present a multi-objective optimisation framework for automatic calibration of Cellular Automata land-use models with multiple dynamic land-use classes. The framework considers objectives related to locational agreement and landscape pattern structure, as well as the inherent stochasticity of land-use models. The framework was tested on the Randstad region in the Netherlands, identifying 77 model parameter sets that generated a Pareto front of optimal trade-off solutions between the objectives. A selection of these parameter sets was assessed further based on heuristic knowledge, evaluating the simulated output maps and parameter values to determine a final calibrated model. This research demonstrates that heuristic knowledge complements the evaluation of land-use models calibrated using formal optimisation methods.

## 2.1 Introduction

Modelling of land-use change plays an important role in many areas of environmental planning, such as river basin management (Van Delden et al., 2007), natural area preservation (Hewitt et al., 2014), the development of sustainable agricultural practises (Murray-Rust et al., 2014a, Murray-Rust et al., 2014b), and the influence of urban dynamics on surrounding regions (Haase et al., 2012, Lauf et al., 2012). To better understand the influences of land-use changes, models are increasingly being used as part of decision support systems to evaluate policy that influences spatial planning (Van Delden et al., 2011). To represent land-use dynamics realistically, such models must incorporate complex socio-economic and biophysical drivers with human-environment interactions (Lambin et al., 2001). As a result, Land-Use Cellular Automata (LUCA) have become a popular modelling framework for evaluating land-use changes, as they are able to simulate the behaviour of complex systems with a high degree of realism (Hewitt et al., 2014).

Historically, Cellular Automata methods were proposed for application to geographic systems by Tobler (1979), with LUCA models first used to replicate observed fractal patterns of urban evolution (Couclelis, 1985, Couclelis, 1989, Batty and Longley, 1994), followed by their development into dynamic land-use models (White and Engelen, 1993c, Clarke et al., 1997). Much effort has been invested in developing LUCA models for different global regions, with applications reviewed by Santé et al. (2010). This includes the advent of generic spatial modelling platforms SLEUTH (Clarke et al., 1997) and Metronamica (Van Delden and Hurkens, 2011), which provide well tested models for a range of applications to different study regions. With such generic platforms simplifying model development requirements significantly, research focus on the calibration of LUCA models has increased in recent years (e.g. Blečić et al., 2015, Cao et al. 2014, Garcia et al., 2013, Li et al., 2013, Van Vliet et al., 2013b, Van Vliet et al., 2016).

Calibration of a land-use change model is the process of determining a model parameter set, through the initial setting of model parameters, the iterative adjustment of these parameters based on comparison of the model output with observations, and the selection of a final parameter set, for application to a specific case for long term scenario analysis (adapted from Van Vliet et al. (2016)). The iterative adjustment stage of calibration of LUCA models is extremely complex, as land-use change is a path dependent process that is driven by multiple interdependent processes with uncertain outcomes (Brown et al., 2005). Conventionally this stage of calibration of LUCA models is manual (Van Delden et al., 2012), incorporating the

modeller's process understanding to address this inherent uncertainty. However, implementation of such methods is time consuming, subjective (Jafarnejhad et al., 2016), and lacks transparency and repeatability (García et al., 2013). Consequently, in order to make parameter adjustment more efficient and repeatable, there has been an increasing focus on automating this process (Van Vliet et al., 2013b).

Automatic parameter adjustment methods generally make use of formal optimisation methods that maximise model performance metrics (Blecic et al., 2015, Cao et al., 2014, García et al., 2013, Li et al., 2013). Consequently, the success of these methods relies heavily on the ability to assess performance in a quantitative fashion. This assessment has to consider two separate properties of LUCA model performance: (i) locational agreement, alternatively termed cell-by-cell agreement (Hagen-Zanker, 2009), which is the match of pixels between simulated outputs and the corresponding observed data (Van Vliet et al., 2013a, Hagen-Zanker, 2009), and (ii) landscape pattern structure, which is the inferred realism of land-use change processes captured by the difference between observed and simulated landscape patterns (Engelen and White, 2008). Consequently, automatic parameter adjustment of LUCA models can be considered a multi-objective optimisation problem (Hagen-Zanker, 2008).

At present, multi-objective optimisation has only been applied to the parameter adjustment stage of calibration of LUCA urban growth models that are implemented with two land-use classes, despite the capacity of these models to consider a broader range of land-use classes, using the SLEUTH metrics (Trunfio, 2006) or logit regression model fitness functions (Cao et al., 2014). Whilst this work has merit in characterising urban and non-urban interactions, it represents a less complex calibration problem than LUCA models that consider multiple dynamic land-use classes, as these are more complex models that possess a significantly larger number of parameters for calibration. In contrast, studies that have used optimisation approaches for the calibration of LUCA models with multiple dynamic land-use classes have only considered a single objective. For example, Blecic et al. (2015) considered only the locational agreement element of performance for parameter tuning, and while García et al. (2013) used both locational agreement and landscape pattern structure metrics, these were combined into a single objective during the optimisation process, where as a result, important trade-offs between locational agreement and landscape pattern structure could not be examined. Both studies also generated only one possible model parameterisation for future scenario analysis, limiting the ability to understand how calibrated parameters are potentially influenced by the metrics used for optimisation.



To address these shortcomings, the objectives of this paper are (i) to present a multi-objective optimisation framework that automates the parameter adjustment stage of calibration of LUCA models with multiple dynamic land-use classes, enabling the identification of multiple model parameter sets that could be suitable for long-term scenario analysis; and (ii) to demonstrate the application of the framework on the case study comprising the Randstad region in the Netherlands. The remainder of this paper is organised as follows: The proposed multi-objective optimisation-based calibration framework is introduced in Section 2.2, followed by a description of an application to a case-study of Randstad in Section 2.3. The results for the case study are presented and discussed in Sections 2.4 and 2.5. The conclusions and recommendations of this work are presented in Section 2.6.

## **2.2 Proposed multi-objective optimisation-based calibration framework**

The proposed multi-objective optimisation framework for calibration of LUCA models with multiple dynamic land-use classes is presented in Figure 2.1. As shown, the framework is comprised of four stages. First, in the selection stage, the components required for optimisation are chosen. Next, in the specification stage, to ensure an efficient and robust output, certain aspects prevalent to the previously selected components are specified. Following this, the multi-objective optimisation parameter adjustment is implemented and run to completion. Finally, the resulting model outputs are assessed, quantitatively evaluated using a neutral model, followed by heuristic interpretation of the outputs to decide on a final model parameter set.

### ***2.2.1 Selection stage***

In the selection stage, the four main components for optimisation are chosen, as shown in Figure 2.1: the LUCA model to be used and the parameters to be adjusted, the optimisation algorithm used for the parameter adjustment process, and the map comparison metrics used to assess model performance.

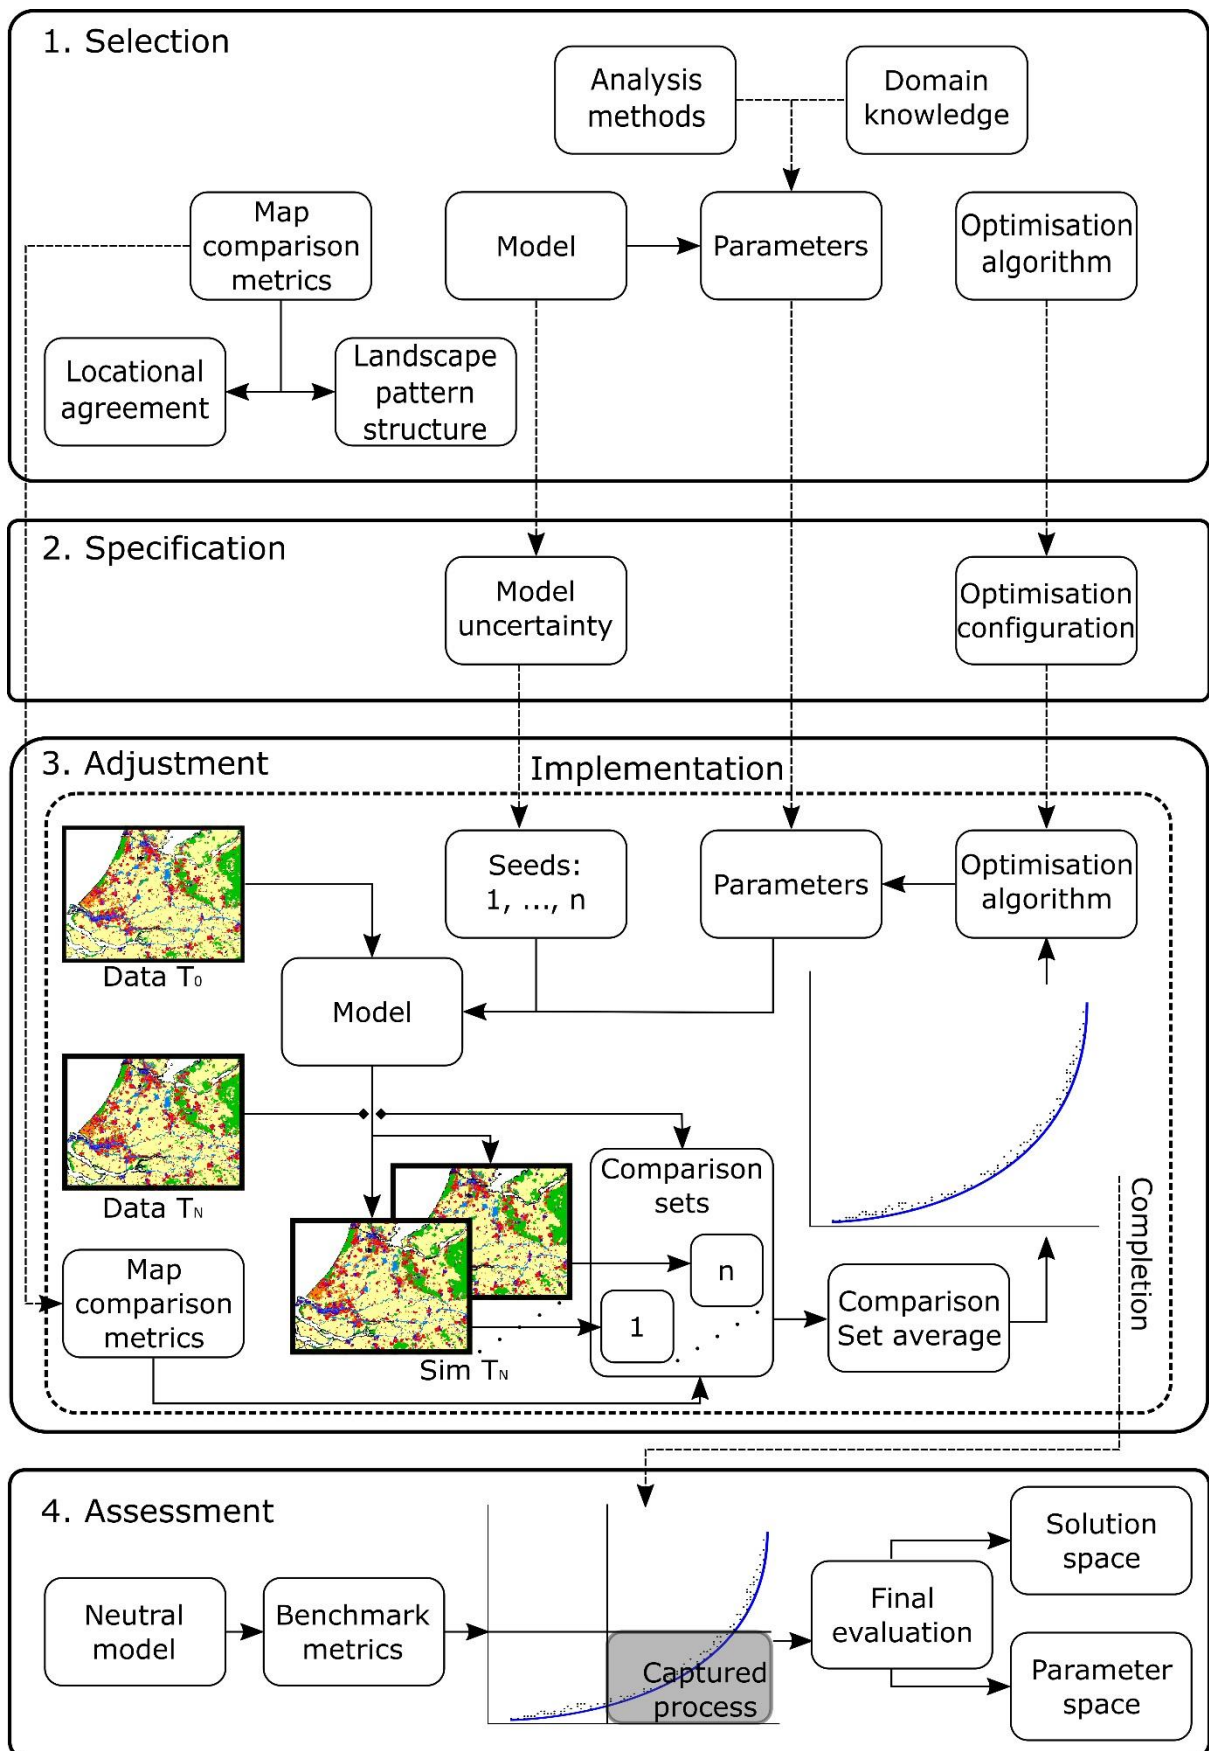


Figure 2.1. Proposed multi-objective optimization-based calibration framework for LUCA model with multiple dynamic land-use classes

### *2.2.1.1 Model and parameters*

The LUCA model determines the number and type of parameters that require adjustment. The parameters within LUCA models are used to capture the processes that influence land-use changes, such as the physical suitability of the landscape and the influence different land-use classes exert on each other. The consideration of multiple dynamic land-use classes, as is the case for transition potential models derived from White and Engelen (1993c), introduces a large number of parameters that must be adjusted, to capture the respective influences of each process on different land-use classes (García et al., 2013). The most notable example of parameters for transition potential models are neighbourhood rules, which characterise the influence different land-use classes exert on each other at different distances (RIKS, 2015). For example, considering a neighbourhood size of eight cells introduces 30 parameters for each neighbourhood rule. As the number of neighbourhood rules is the product of the total number of land-use classes and the number of actively allocated land-use classes, the resulting number of parameters that could be adjusted is large. Consequently, it is desirable to be judicious about which parameters to include in the automatic adjustment process.

There are several approaches for selecting which parameters to include in the automatic adjustment process, including empirical understanding of the region of interest, in order to select the parameters that correspond to the dominant processes driving land-use change (Hewitt et al., 2014) and the use of quantitative analysis methods, such as data mining or sensitivity analysis (Li and Yeh, 2004, Wang et al., 2011b, Gibbs et al., 2012). Alternatively, the computational burden associated with the automatic adjustment process can be reduced by using empirical methods to identify reasonable values of certain parameters, such as neighbourhood rules (Van Vliet et al., 2013b), which can then be used to initialise the optimisation at a good solution, which has been shown to reduce the computational demands of optimisation problems in other fields (e.g. Bi et al., 2016).

In addition to the number of parameters that require adjustment, the size of the space of potential parameters values that need to be explored during the automatic adjustment process, which is commonly referred to as the search space, is also affected by the potential ranges different parameters can take (Maier et al., 2014). Consequently, the upper and lower limits of the parameters that are included in the automatic calibration process need to be selected carefully, balancing the needs to ensure the best possible combination of parameter values can be identified with the desire to reduce the size of the search space. The size of the search space can also be reduced by restricting the values parameters can take to be in accordance with

underlying domain knowledge. For example, neighbourhood effects can be parameterised to generate shapes consistent with discursive knowledge (Hagoort et al., 2008), reducing the number of parameters from the order of thousands to hundreds (Blecic et al., 2015, García et al., 2013).

#### *2.2.1.2 Multi-objective optimisation algorithm*

As part of the proposed framework, it is suggested a population-based metaheuristic algorithm is used for the automatic parameter adjustment process. Such algorithms are advantageous because they are able to find (near) globally optimal solutions for highly complex (e.g. non-linear, non-convex) problems, and it is straightforward to link them with existing simulation models, such as LUCA models, without the need for problem simplification (Maier et al., 2015). This is because they work in an iterative fashion using a population of separate solutions, where the model parameters are adjusted for each member of the population based on the search strategy of the algorithm under consideration (e.g. survival of the fittest in the case of genetic algorithms or the foraging behaviour of ants in their search for food in the case of ant colony optimisation), and how well these adjustments perform is evaluated by running the model with the altered parameters. Information on how well the suggested changes have performed are fed back to the algorithm, informing which changes are made to the model parameters in the next iteration (or generation) and so on. In addition, metaheuristic algorithms have the capacity to handle multiple objectives (Maier et al., 2014), enabling them to optimise measures of locational agreement and landscape pattern structure simultaneously.

There are a number of population-based metaheuristics (Kingston et al., 2008), which can be used to automate LUCA model parameter adjustment, utilising various heuristic mechanisms, with heuristic selection mainly due to preference. Applied algorithms include ant colony optimisation (Liu et al., 2012), memetic algorithms (Veerbeek et al., 2015), particle swarm algorithms (Blecic et al., 2015, Feng et al., 2011) and genetic algorithms (Cao et al., 2014, García et al., 2013, Li et al., 2013, Trunfio, 2006, Clarke-Lauer and Clarke, 2011). The proposed framework caters to different population-based metaheuristics, but requires an algorithm that can handle multiple objectives, to allow for exploration of the trade-offs between the objectives.

#### *2.2.1.3 Map comparison metrics*

As mentioned previously, optimisation requires quantifiable objectives to assess performance. Consequently, the proposed framework requires the selection of metrics that can quantify the

objectives of maximising locational agreement and minimising landscape pattern structure error. For locational agreement, potential metrics include *percentage correct* (Santé et al., 2010), the *figure of merit* (Pontius Jr. et al., 2008), *allocation and quantity agreement* (Pontius Jr. and Petrova, 2010), or *Cohen's Kappa* or one of its variations (Van Vliet et al., 2011, Van Vliet et al., 2013b). There are also many landscape pattern structure metrics (McGarigal, 2014) that quantify different aspects of the landscape and have been applied previously in land-use modelling studies. These include (amongst others) *edge density* (García et al., 2013), *largest patch index* (Li et al., 2013), and *clumpiness* (Van Delden et al., 2012).

### 2.2.2 *Specification stage*

As shown in Figure 2.1, this stage follows the selection stage, and requires the specification of certain relevant components of the model, and optimisation algorithm, to allow for effective optimisation. This includes how model stochasticity is taken into account, and the parameterisation of the selected optimisation algorithm.

#### 2.2.2.1 *Model stochasticity*

LUCA models generally include a stochastic element to capture the variability of human decisions that drive land-use changes, by including a random perturbation factor in the model. This stochasticity must be considered appropriately to gain meaningful optimisation results. In order to achieve this, Guo et al. (2017) recommend using a number of model runs with different random seeds, but to use the same seeds during each iteration of the optimisation in order to ensure the impact of changing model parameters from one optimisation iteration to the next is not diluted or confused by the stochastic nature of the model used to assess the objective function values. Therefore, as part of the proposed framework,  $n$  LUCA model runs with different random number seeds are used to assess model performance for a given set of model parameters during each iteration of the automatic parameter adjustment optimisation process. However, the same random number seeds are used in every iteration of the optimisation process to ensure that any changes in the objectives from one iteration to the next are due to changes in model parameters, rather than a combination of these changes and randomness in LUCA models.

#### 2.2.2.2 *Optimisation configuration*

Three aspects of the optimisation process must be specified prior to running the optimisation algorithm. First, values of the parameters that influence optimisation algorithm searching behaviour, such as population size, probability of cross-over and probability of mutation in the

case of genetic algorithms, must be defined (Zheng et al., 2016, Zecchin et al., 2012). It should be noted that these parameters are different from the parameters of the LUCA models that are to be determined with the proposed automatic parameter adjustment process. Second, termination criteria for the optimisation process, such as a pre-defined number of iterations or no significant improvement in performance, must be specified. Finally, the number of times the entire optimisation process has to be repeated,  $R$ , has to be specified. Such repetition is needed due to the stochastic nature of population-based metaheuristics, and increases the chance that the best possible combination of LUCA model parameters is identified and that the results of the automatic parameter adjustment process are robust. It should be noted that the stochastic nature of the optimisation process is distinct from the uncertainty associated with the LUCA models. While the LUCA model is run  $n$  times during each iteration of a single optimisation run in order to account for the stochasticity of LUCA models, the entire optimisation process is repeated  $R$  times to account for the stochasticity of the metaheuristic optimisation process.

### **2.2.3 Adjustment stage**

#### *2.2.3.1 Optimisation process*

The purpose of this stage of the proposed framework is to use the selected optimisation algorithm (see Section 2.2.1.2) to identify the combinations of LUCA model parameters (see Section 2.2.1.1) that provide the best possible trade-offs between the selected performance metrics (see Section 2.2.1.3). As mentioned previously, this is achieved in an iterative fashion by using a metaheuristic optimisation algorithm, to determine changes to the model parameters at each generation of the optimisation process based on resulting changes in the performance metrics. The steps in this process are shown schematically in Figure 2.1, and in more detail in Figure 2.2.

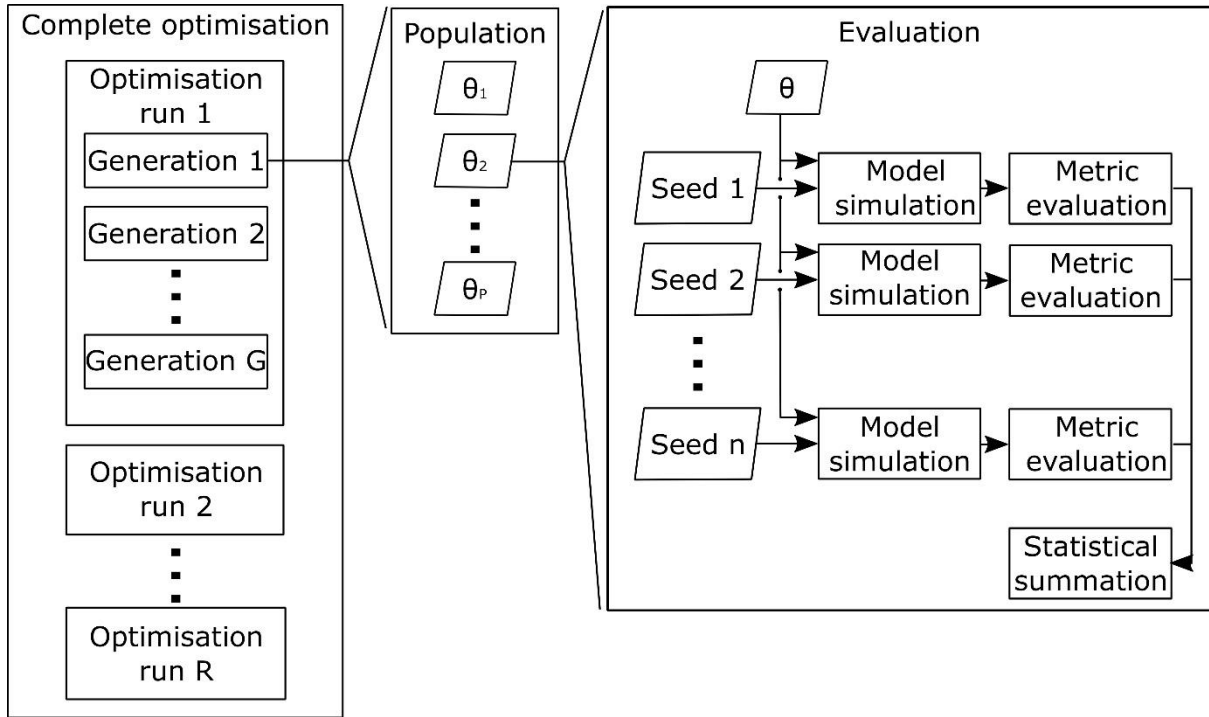


Figure 2.2. Population-based metaheuristic for optimisation-based LUCA model parameter adjustment process

As can be seen in Figure 2.2, given that population-based metaheuristics are used for optimisation, a population of  $P$  solutions ( $P$  sets of LUCA model parameter values) is generated for each iteration of the optimisation process. Next, a LUCA model with each of these  $P$  sets of parameter values  $\theta$  (the parameters selected for optimisation input into the model) is run from time  $T_0$  until time  $T_N$ , with an actual land-use map at time  $T_0$  providing the initial conditions for each of the  $P$  model runs (Figure 2.1). Each simulated output map at time  $T_N$  is quantitatively compared with a corresponding map of actual land-use using the metrics in the comparison set selected previously; one to assess locational agreement and another to assess the error of landscape pattern structure. As mentioned previously, given the stochastic nature of land-use change, each of the  $P$  land-use models is run with  $n$  different random number seeds, generating  $n$  simulated land-use maps, each of which is individually compared with the corresponding map of actual land-use. The performance of each model with the  $P$  sets of selected model parameters is quantified by taking the average of the metric values across the  $n$  different outputs (Figures 2.1 and 2.2). Hence, the final metric value for one member of the population of the metaheuristic is given by the following objective:

$$LA_p(\theta) = \frac{1}{n} \sum_{i=1}^n LA_i(\theta) \quad (2.1)$$

where  $LA_p$  is the locational agreement objective function value of the member of the population  $p$ ,  $\theta$  are the input parameters corresponding to that member of the population, and  $n$  is the number of LUCA model random number seeds being considered. Similarly, for landscape pattern structure:

$$LPS_p(\theta) = \frac{1}{n} \sum_{i=1}^n LPS_{Error,i}(\theta) \quad (2.2)$$

where  $LPS_p$  is the landscape pattern structure error objective function value of the member of the population  $p$ , and  $n$  is the number of LUCA model random number seeds.

As mentioned previously, when metaheuristic optimisation algorithms are used,  $\theta$  values are optimised in an iterative fashion with the aim of improving the performance metrics from one generation to the next. Consequently, the above steps are repeated for a total of  $G$  generations (Figure 2.2) until the desired stopping criteria have been met (see Section 2.2.2.2). Finally, as also mentioned previously, each optimisation run is repeated  $R$  times to account for the stochastic nature of the searching behaviour of population-based metaheuristics.

As parameter adjustment is a multi-objective optimisation problem, the optimisation process does not generate a single set of model parameters. Rather, the optimisation generates a series of model parameter sets that generate optimal, non-dominated trade-offs between the two objectives. The output of the optimisation is shown in Figure 2.1 by the curve called the, “Pareto front of objectives.” The Pareto front indicates the optimal trade-off between objectives, where improved performance in one objective cannot be achieved without inferior performance in the other objective.

### 2.2.3.2 Implementation

The proposed optimisation process is very computationally expensive, because of the relatively long simulation times associated with running LUCA models (e.g. run times can vary from 10 seconds to 10 minutes, depending on the number of classes, spatial extent and resolution) and the large number of simulations required, as illustrated in Figure 2.2. For example, for typical values of  $P=200$ ,  $n=10$ ,  $G=500$  and  $R=5$ , the LUCA model would have to be run  $200 \cdot 10 \cdot 500 \cdot 5=5,000,000$  times, resulting in a total run time ranging from 50,000,000 to 3,000,000,000 seconds, or 600 to 35,000 days, for individual LUCA model run times of 10 and 600 seconds, respectively, which is not feasible from a practical perspective. However, as has been discussed



in other problem domains, there are a number of avenues for increasing the computational efficiency of this process, including the incorporation of heuristic information to reduce the size of the search space (Szemis et al., 2012, and Section 2.1.1) or to improve the efficiency of the optimisation process (Nguyen et al., 2016), by using meta-models as surrogates for the computationally expensive LUCA models (Broad et al., 2015) or the use of parallel computing resources (Blecic et al., 2015). By using these approaches, automatic parameter adjustment methods can be made feasible for application to complex LUCA models.

#### **2.2.4 Assessment stage**

To determine the final calibrated model, the resultant model parameter sets corresponding to different Pareto front solutions are assessed using a comprehensive set of methods that have been used in similar calibration studies (Van Vliet et al., 2013b, Hewitt et al., 2014) including assessment of objective output using a neutral model benchmark test, and a final evaluation comprising heuristic interpretation of the solutions and parameter values obtained.

##### *2.2.4.1 Neutral model benchmark test*

An issue with calibrating LUCA models is the impact of boundary conditions. Due to land-use commonly persisting, metrics often indicate good model performance, though this can largely be due to limited variation of the final landscape from the initial situation (Hagen-Zanker and Lajoie, 2008). Hence, the proposed framework uses a neutral model benchmark test to evaluate the output in the objectives space, to determine which model parameter sets have captured land-use change processes appropriately. This is illustrated by the black lines over-laid on the Pareto front in the assessment stage in Figure 2.1. A model that outperforms the benchmark (i.e. generates superior metric values) is considered valid, because model performance can be attributed to correct capture of processes, rather than inherent land-use persistence (Hagen-Zanker and Lajoie, 2008).

##### *2.2.4.2 Final evaluation*

A final evaluation of the optimisation results is performed by assessing the solution and parameter spaces (Figure 2.1). This can be performed by a modelling expert, or in a participatory manner, where stake-holders are included in the evaluation to decide on the most appropriate model, as with manual calibration (Hewitt et al., 2014). Evaluation of the solution space refers to interpreting the land-use maps generated by the different model parameter sets, corresponding to different solutions along the Pareto front. Evaluation of the parameter space refers to assessing the parameters corresponding to the different solutions, to determine if these

are consistent with expected land-use change processes of the study region. By doing so, a final model parameter set can be recommended.

### 2.3 Case study application of proposed framework

The proposed framework outlined in Section 2.2 is applied to a real-world case-study of the Randstad region in the Netherlands from 1989 to 2000. Randstad is a conurbation of the four largest cities in the Netherlands and the surrounding region. The model for the case study region has a 500-metre resolution and covers a spatial extent of 14,175 km<sup>2</sup>, shown in Figure 2.3.

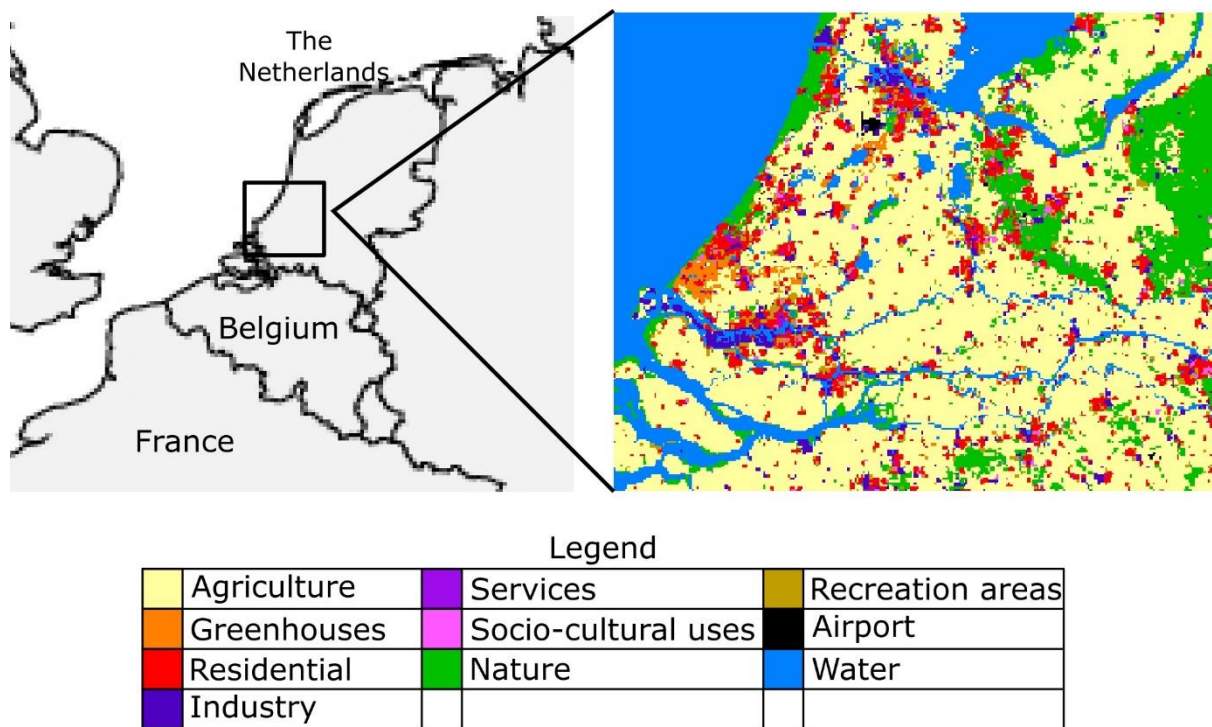


Figure 2.3. Location of the Randstad region and rasterized map for 1989

#### 2.3.1 Selection stage

##### 2.3.1.1 Land-use model

The transition potential LUCA model Metronamica is used (Van Delden and Hurkens, 2011). Metronamica is a generic modelling platform, evolved from the pioneer model by White and Engelen (1993c) that facilitates direct application to the study region of interest, capturing the dynamics between urban and regional systems (RIKS, 2015), which is able to consider multiple dynamic land-use classes. It has been used successfully in numerous decision support systems in diverse global regions, such as the regional and national socio-economic policies in the Netherlands (Engelen et al., 2003), long-term regional planning in Waikato, New Zealand (Rutledge et al., 2008), river basin management for Mediterranean watersheds (Van Delden et

al., 2007), assessment of the impact of intra-urban land prices (Furtado et al., 2012), impact assessment of agricultural policies in Europe (Van Delden et al., 2010) and modelling shifting cultivation practises in Sri Lanka (Wickramasuriya et al., 2009).

In Metronamica, a transition potential is used to allocate land-use classes, calculated to update land-use at each time step, using various factors to define the potential for each cell to support each land-use class considered. Land-use classes are allocated until all demand requirements are met (demands for particular land-uses are exogenously defined), based on the transition potential for each given class, or until there are no locations available. Metronamica considers three land-use class categories: active, which are dynamically modelled and allocated (commonly urban classes); features, which are fixed land-uses that influence landscape dynamics but do not have dynamic behaviour (such as airports or water bodies); and passive, which occupy the remaining landscape and change in extent as a result of changes to the other categories. These are normally occupied with land-use types that have low transition costs (for example agricultural land). The Metronamica model developed for the Randstad case-study has ten land-use classes. There are seven active, *Greenhouses, Residential, Industrial, Services, Socio-cultural uses, Nature and Recreation areas*; one passive, *Agriculture*; and two features, *Airport and Water*.

To determine the transition potential, Metronamica considers four processes: accessibility, the provision of infrastructure required for certain land-use classes; suitability, the influence of physical factors; neighbourhood effects, the spatial interactions between land-use classes representing the behaviour of actors and activities taking place on those classes; and zoning, the influence of spatial planning. A stochastic component is also included to incorporate the inherent uncertainty of land-use change decisions. The transition potential is calculated by:

$$TP_{c,k} = A_{c,k} \cdot S_{c,k} \cdot N_{c,k} \cdot Z_{c,k} \cdot \gamma \quad (2.3)$$

where  $TP_{c,k}$  is the transition potential for the cell  $c$  to support land-use class  $k$  (subscripts for all symbols above have the same interpretation),  $A_{c,k}$  is the accessibility,  $S_{c,k}$  is the suitability,  $N_{c,k}$  is the neighbourhood effect,  $Z_{c,k}$  is the zoning influence, and  $\gamma$  is the stochastic component, the form of which is given in the Metronamica documentation (RIKS, 2015). This documentation also includes further details of the transition potential and underlying equations for each process. The Metronamica model developed for Randstad includes two accessibility layers; motorways and other roads. Four physical attributes are included for suitability; noise,

elevation, natural hydrology and slope. Zoning for ecological corridors, the Noise contour for Schiphol and urban expansion plans are also included.

### 2.3.1.2 Parameters

As the case study is a region of urban growth with a short time span, the main sources of land-use conversion will relate to the expansion of the existing socio-economic land-uses. Hence, the major processes driving land-use changes will be the self-organising behaviour of the system for increased expansion of the urban cores (Couclelis, 1989, Batty and Longley, 1994, White and Engelen, 1993a, White and Engelen, 1993b). As a result, the most important processes to include for tuning are those driving this behaviour, the neighbourhood interactions and accessibility (Verburg et al., 2004). Consequently, parameters for neighbourhood rules and accessibility are included for automatic tuning.

Neighbourhood rules are parameterised to define the influence that different land-use classes exert on each other. The cumulative transition potential due to neighbourhood influence for the conversion from one land-use to another is calculated by:

$$N_{c,k} = \sum_{c' \in D(c)} w_{k,K(c'),d(c,c')} \quad (2.4)$$

where  $N_{c,k}$  is the neighbourhood transition potential for cell  $c$  transitioning to land-use type  $k$ ,  $D(c)$  is the set of all cells in the neighbourhood of the cell of interest  $c$ ,  $K(c')$  is a look-up function that returns the land-use class for cell  $c'$ ,  $d(c,c')$  is the distance between cells  $c$  and  $c'$ , and  $w_{k,j(x)}$  expresses the influence that a cell with land-use class  $j$ , returned by the look-up function  $K(c')$ , exerts on a cell of potential land-use class  $k$  at a linear distance of  $x$  between the two cells.

Accessibility defines the importance of infrastructure elements for different land-use classes, parameterised as:

$$A_{c,k} = \begin{cases} \frac{a_{s,k}}{D_{s,c} + a_{s,k}} & \text{if } a_{s,k} > 0 \\ 0 & \text{if } a_{s,k} = 0 \\ 1 - \frac{|a_{s,k}|}{D_{s,c} + |a_{s,k}|} & \text{otherwise} \end{cases} \quad (2.5)$$

where  $D_{s,c}$  is the cellular distance between cell  $c$  and the nearest cell containing the type of infrastructure  $s$ , and  $a_{s,k}$  is the accessibility decay parameter that expresses the importance of the type of infrastructure  $s$  to land-use  $k$ . Values for  $a_{s,k}$ , and a weighting parameter are tuned

for each actively allocated land-use class  $k$ , and each infrastructure type  $s$ . For this case-study 28 parameters are tuned for accessibility, because there are 7 actively allocated land-use classes, and 2 types of infrastructure layers (motorways and other roads).

To reduce the size of the search space during the automatic parameter adjustment stage, the neighbourhood rules are delineated into two specific parts, as shown in Figure 2.4: the locus point (dot), which defines the inertia (self-influential rules) or conversion (interactive rules) of the location of interest; and the tail (line), which defines the interaction effects for different distances. There are several common shapes that the neighbourhood rule tails can take (Van Vliet et al., 2013b, Hagoort et al., 2008). The most common forms are characterised by a high influence at shorter distances, a point of inflection, and a slow, gradually decaying influence over large distances. To capture such dynamics, this research parameterises neighbourhood effect tails using exponential decay functions, thereby reducing the number of parameters that require calibration, of the form:

$$y(x) = \begin{cases} c & \text{for } x = 0 \\ ae^{-bx} & \text{for } 0 < x \leq x_c \\ 0 & \text{for } x > x_c \end{cases} \quad (2.6)$$

where  $a$  and  $b$  are the controlling parameters of the neighbourhood rule,  $x$  is the distance,  $y(x)$  is the influence value,  $c$  is the locus point of persistence and conversion, and  $x_c$  is the critical distance where the influence is set to zero. For each neighbourhood rule, three parameters are tuned,  $a$ ,  $b$  and  $c$ . For this case-study, there are 10 land-use classes and 7 dynamic land-use classes, hence 70 neighbourhood rules are included for adjustment. By parameterising neighbourhood rules using exponential decay functions, 210 parameters are tuned for neighbourhood rules. Hence, a total of 238 parameters are tuned as part of the adjustment process for this case study.




Figure 2.4. Parameterisation of a neighbourhood effect using exponential decay function, a locus point and critical distance

### 2.3.1.3 Multi-objective optimisation algorithm

This research uses the Non-dominated Sorting Genetic Algorithm II (NSGA-II) first proposed by Deb et al. (2002), which is a population-based Multi-Objective Genetic Algorithm (MOGA) regarded as an industry standard (Wang et al., 2015) with a demonstrated ability to tune LUCA model parameters (Cao et al., 2014, Trunfio, 2006). The MOGA is used to find multiple alternative LUCA parameter sets that represent the best possible trade-off between the calibration objectives, the Pareto front of solutions, illustrated in Figure 2.1 by the curved blue line.

MOGAs use simple heuristics (computationally efficient rules) across a number of generations, to derive information about which decision variable values result in better performing objectives, and use this information about the present (parent) population to generate the next (child) generation of solutions, some of which are likely to have superior objective function values. The heuristics mimic those of, ‘survival of the fittest,’ and are referred to as selection, cross-over and mutation. Selection is used to promote better performing solutions to the subsequent generations, by comparing different solutions and selecting, by some mechanism, those which perform better. Cross-over takes a subset of decision variable values from a pair of parent solutions and randomly recombines them to form new child solutions. The purpose is to exploit good solutions, as better performing decision variable values from the parents, when combined differently, may result in children with superior performance. This is governed by the probability of cross-over, which is set prior to optimisation. Mutation takes a small

subset of decision variable values in a child and perturbs them, the purpose of which is to diversify the search to explore a wider possible range of solutions, which might lead to superior objective function values.

The specific MOGA operators employed by the NSGA-II, as well as the additional features that distinguish it from conventional MOGAs, are illustrated in Figure 2.5. As shown, NSGA-II begins with an initialisation step, where an initial population of solutions with randomly generated decision variable values is created. The associated objective function values are then calculated. Next, the standard genetic operators previously discussed (selection, cross-over, mutation) are used to generate a child population from the existing parent population. This implementation of NSGA-II uses tournament selection, simulated binary cross-over, and polynomial mutation, as implemented by Deb et al. (2002), because this problem uses real-value decision variables.

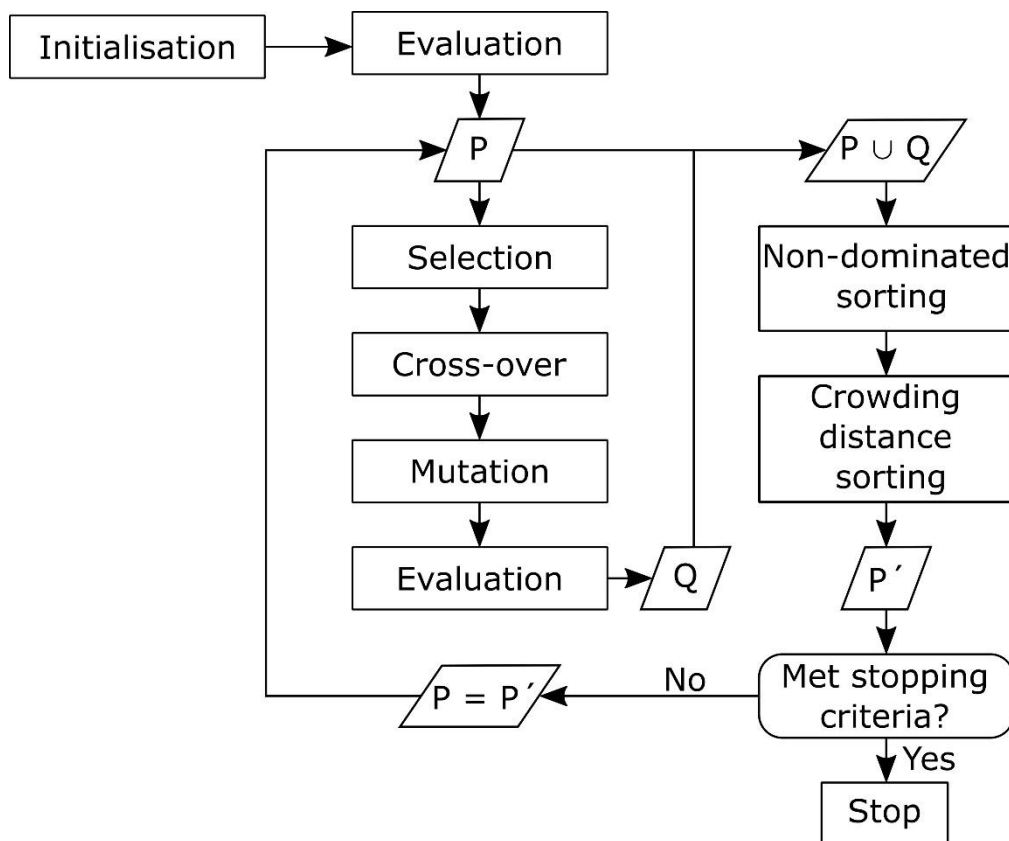


Figure 2.5. Optimisation process performed by NSGA-II, the parent population ( $P$ ) is modified to generate a child population ( $Q$ ), then recombined and modified further to produce the parent population ( $P'$ ) of the next generation

The key performance advantage of using NSGA-II is achieved via the additional operators implemented in the algorithm, which preserve non-dominated (i.e. the best performing) solutions from generation to generation. As shown in Figure 2.5, NSGA-II recombines the parent ( $P$ ) and child ( $Q$ ) populations, and ranks the combined population in order of non-dominated fronts that occur, which are a series of Pareto fronts. From this, a new parent population ( $P'$ ) is formed, by iteratively selecting the solutions belonging to the best fronts. If only a subset of solutions from a front can be included, crowding distance sorting is used, which aims to select the most diverse set of solutions from the final front that is included. This process is repeated until the stopping criteria are met.

#### 2.3.1.4 Map comparison metrics

For the case study, locational agreement is measured using Fuzzy Kappa Simulation (FKS) developed by Van Vliet et al. (2013a). The FKS metric is an adaptation of Cohen's Kappa (Santé et al., 2010, Van Vliet et al., 2016), which measures the fuzzy agreement of location and class between two data sets relative to a baseline of random allocation:

$$FKS(\theta) = \frac{PO(\theta) - PE(\theta)}{1 - PE(\theta)} \quad (2.7)$$

where  $PO(\theta)$  is the observed agreement, the similarity of the simulated land-use transitions with the observed land-use transitions, and  $PE(\theta)$  is the expected agreement, the agreement obtained from a random allocation of the given class transitions relative to the initial land-use map. FKS is used because it measures the agreement of land-use transitions, and hence includes an implicit baseline. It also uses fuzzy interpretation, which attributes partial agreement relative to proximity. In this research, partial locational agreement of correct transitions is considered for a neighbourhood radius of two cells, with strength of agreement using a halving distance of one. Hence, the locational agreement measure for a single member of the population ( $LA_p(\theta)$ ) is equivalent to FKS agreement, averaged across the total number of random number seeds considered.

To assess the error of landscape pattern structure, the average of the absolute class level clumpiness error between actual and simulated values is used. This metric is used based on successful previous applications as a part of the manual parameter tuning of LUCA models (Van Delden et al., 2012). Clumpiness is an aggregation metric that measures the proportional deviation of the proportion of like adjacencies involving the corresponding class from what is expected under a spatially random distribution (McGarigal, 2014). It is calculated first by determining the proportion of like adjacencies:



$$G_k = \frac{g_{kk}}{\sum_{l=1}^L g_{kl} - \min e_k} \quad (2.8)$$

where  $G_k$  is the proportion of like adjacencies for land-use class  $k$ ,  $g_{kl}$  is the count of like adjacencies between patches of class  $k$  and  $l$  using the double-count method (McGarigal, 2014),  $L$  is the total number of land-use classes,  $\min e_k$  is minimum perimeter of a maximally clumped patch of class  $k$ , defined as:

$$\min e_k = \begin{cases} 4n & n^2 = a_k \\ 4n + 2 & n^2 < a_k \leq n(1 + n) \\ 4n + 4 & a_k > n(1 + n) \end{cases} \quad (2.9)$$

where  $a_k$  is the area of class  $k$  (in terms of number of cells), and  $n$  is the length of a side of the largest integer square possible with a smaller area than  $a_k$ . With the proportion of like adjacencies calculated, it follows that clumpiness is calculated by:

$$CLUMPY_k = \begin{cases} \frac{G_k - P_k}{P_k} & \text{for } G_k < P_k \text{ and } P_k < 0.5 \\ \frac{G_k - P_k}{1 - P_k} & \text{otherwise} \end{cases} \quad (2.10)$$

where  $P_k$  is the proportion of the landscape occupied by patch type (class)  $k$ . As clumpiness is calculated at the class level, it is aggregated by taking the average across the dynamic land-use classes. Hence, for this case study, landscape pattern structure error is given by the average error between the observed and simulated clumpiness values ( $\Delta CLU$ ) for the seven actively allocated land-use classes, averaged across the total number of random number seeds considered.

## 2.3.2 Specification stage

### 2.3.2.1 Model stochasticity

Stochasticity in the Metronamica model is accounted for by calculating the map comparison metrics using ten stochastically generated replicates of simulated land-use maps for each iteration of NSGA-II (i.e.  $n=10$ ) in accordance with Equations 2.1 and 2.2.

### 2.3.2.2 Optimisation configuration

Although it is generally recommended to fine tune the parameters controlling the searching behaviour of the optimisation algorithm, such as the population size, the probability of mutation and probability of cross-over, via sensitivity analysis before application, for large problems, this is extremely computationally expensive and therefore impractical (Wang et al., 2015). As this research falls into this category of problems, recommended NSGA-II parameter values are adopted (Zheng et al., 2016, Wang et al., 2015), namely a probability of crossover of 0.9 and a

probability of mutation equalling the inverse of the number of decision variables, 0.0042 in this case. The selection of population size, which corresponds to how many chromosomes would be included per generation, is also informed by the work of Wang et al. (2015) into optimisation problems with many decision variables. Their results show that the population size for large problems must be greater than the number of problem decision variables to find non-dominated solutions. Hence, this research uses a population size of 256 (i.e.  $P=256$ ).

The stopping criterion adopted is based on solution convergence, and reaching a certain number of generations. The required number of generations used is 500 (i.e.  $G=500$ ). Convergence implies that, for the given optimisation algorithm configuration, negligible improvement is obtained in the objective space with continued optimisation. There are numerous metrics for measuring convergence (Maier et al., 2014). This research uses the hyper-volume metric (Zitzler, 1999), a commonly employed metric of multi-objective performance (Reed et al., 2013, Hadka and Reed, 2012) that captures convergence and diversity of the objective space. The solutions are considered to have converged when there is negligible improvement in the hyper-volume for 50 generations, defined as less than 0.5% improvement in the hyper-volume metric. The results for all optimisation runs show that the use of 500 generations is sufficient for hyper-volume convergence corresponding to a 0.5% change.

Given the computational demands of the framework, five different optimisation seeds are used (i.e.  $R=5$ ). Following the convergence of the five optimisation runs, the Pareto optimal solutions are combined to identify the best overall front, as is commonly done in multi-optimisation studies (e.g. Wu et al., 2012).

### **2.3.3 Adjustment stage**

#### **2.3.3.1 Implementation**

For the proposed case-study, running the optimisation on an Intel i5-3470S CPU core processor, a single evaluation took approximately 6 seconds (i.e. a single run of the Metronamica model). Consequently, extrapolating this out to a complete optimisation run, the total computational time would be approximately  $P \cdot n \cdot G \cdot R \cdot 6 = 256 \cdot 10 \cdot 500 \cdot 5 \cdot 6 = 444.4$  days. Hence, to make the optimisation process feasible from a practical perspective, a parallelised version of NSGA-II is implemented, to increase computational efficiency. This uses the Phoenix cluster, a high-performance computing facility operated by Research Services at the University of Adelaide. The separate optimisation runs are parallelised by distributing the 256 evaluations, one for each member of the population, over 129 CPU processing cores

across eight computational nodes. For the Phoenix cluster, each computational node consists of 2 Xeon E5-2698v3 chipsets each containing 16 cores. To control the parallelisation, a master/slave model of computation is implemented, wherein 128 slave processes running on separate cores are used in parallel to evaluate the objective functions for separate members of the population, and the 129<sup>th</sup> is the master process that coordinates the search and runs the recombination (that is the selection, crossover, mutation, non-dominated sorting and crowding distance operators). Parallelisation is achieved using the Message Passing Interface (MPI) using asynchronous communication to improve algorithm efficiency. Using MPI, the master process passes messages containing a set of decision variable values (a LUCA model parameter set) to each of the slaves. Upon receiving this message, the slave runs the LUCA model, evaluates the calibration objectives, and passes a message back to the master containing these objective values. Through this allocation, near linear speedup is achieved as communication time is orders of magnitude lower than time taken to evaluate the calibration objectives, with NSGA-II completing 500 generations within 48 hours. Code, in C++, for this implementation of the NSGA-II is fully open sourced and available as per the details in the software availability section of this paper.

### ***2.3.4 Assessment stage***

#### ***2.3.4.1 Benchmark test***

The benchmark test is conducted by applying two neutral models, the growing clusters (Van Vliet et al., 2013b) and random constraint match (RIKS, 2010) neutral models, which generate maps of the region according to two different growth strategies. As FKS includes an inherent baseline by only considering transitions (Van Vliet et al., 2013a), a benchmark value of 0 is used for FKS. Neutral model benchmark testing is used to evaluate solutions with respect to the landscape pattern structure error. For consistency, the clumpiness of the active classes (the optimisation objective) from five different growing cluster neutral model outputs is averaged, and the error between this and the observed clumpiness for the year 2000 taken as the benchmark landscape pattern structure metric value. Outputs of the optimisation routine with less error than this value are assumed to have captured model processes adequately, and are considered for the final evaluation.

#### ***2.3.4.2 Final evaluation***

A final evaluation of the model parameter sets is conducted by reviewing the solution maps and model parameters obtained after application of the benchmark test in conjunction with an experienced land-use modeller. Through a collaborative evaluation of these solutions, a final

model parameter set is decided upon based on heuristic evaluation, combining understanding of both the model and study region.

## **2.4 Results**

This section presents and evaluates the results of the multi-objective optimisation framework application to the Randstad case study. This comprises three parts, as shown in Figure 2.1. First, the resultant Pareto front is evaluated, to assess the trade-off found between the metrics and the number of solutions that pass the benchmark test. Second, the solution space is assessed, by heuristic evaluation of the simulated output maps. Third, the parameter space is evaluated, through heuristic interpretation of the trends in the parameter values observed across the Pareto front, as well as the interaction between different parameters. The results are not compared with a more common manual calibration method, due to the complexity and highly subjective nature of such an approach.

### ***2.4.1 Identified Pareto optimal solutions***

The Pareto front comprised of non-dominated solutions from the five optimisation runs is shown in Figure 2.6. In total, 77 Pareto optimal solutions were identified, each corresponding to a model parameter set that generated non-dominated metric values. The output metric values are coloured for each parameter set, with metric values for 10 independent simulations shown for each, illustrating the influence of simulation stochasticity. Red points correspond to parameter sets that generated a low  $\Delta\text{CLU}$  and low FKS agreement, blue points correspond to parameter sets that balance the objectives, and pink points show parameter sets that resulted in a relatively high  $\Delta\text{CLU}$  and a relatively high FKS agreement. The black points show solutions that were analysed in further detail to determine a final model parameter set (see Section 2.4.2).

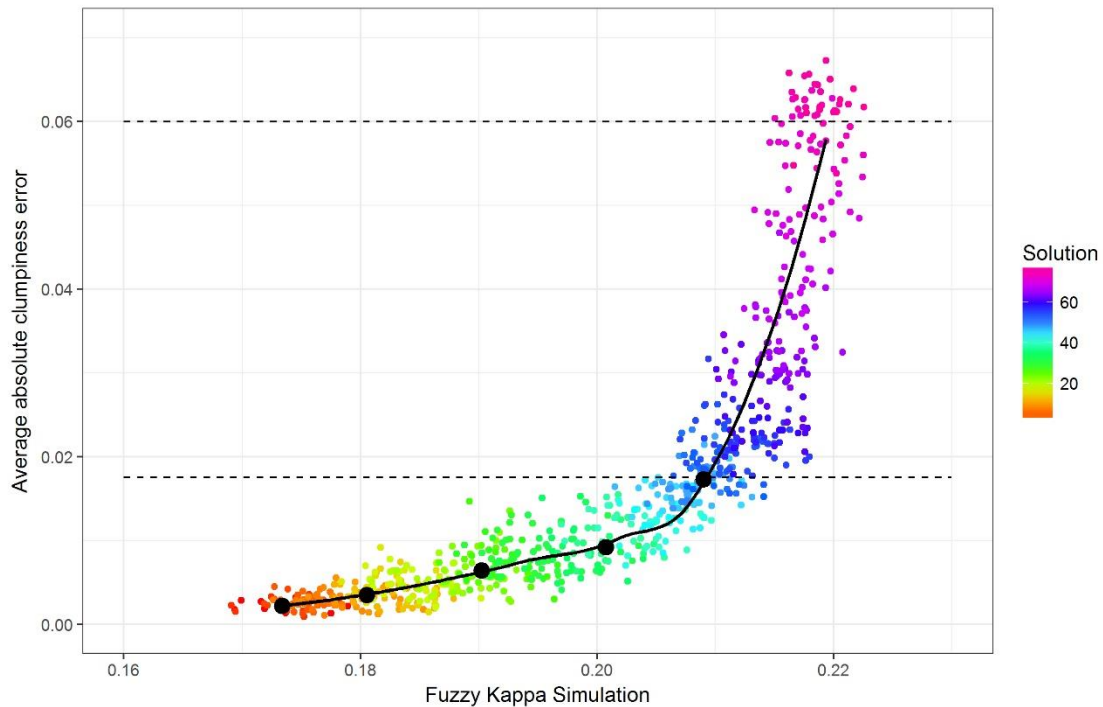


Figure 2.6. Composite Pareto front from the five optimisation runs, consisting of 77 solutions for parameter sets generating non-dominated objectives. The solutions are ordered by ascending FKS value (see colour bar) where the scatter of points arises from 10 independent model runs for each parameter set. The solid black line shows the averaged objective values over the 10 independent model runs. The dashed black lines show the calculated benchmark metric values for the growing clusters neutral model (lower) and random constraint match neutral model (higher). The five black dots show the solutions that are analysed in detail.

As shown in Figure 2.6, there was a meaningful trade-off between the two objectives used. Values of the average absolute clumpiness error across the actively allocated land-use classes varied from 0.001 to 0.067, and values of FKS ranged from approximately 0.169 to 0.222, with a reasonably uniform distribution of points in-between the extremes. The solutions generated corresponded with (the more commonly used) Kappa values between 0.758 and 0.789, comparable with other calibration studies (Cao et al., 2015, Chaudhuri and Clarke, 2013a, García et al., 2013, Jafarnejhad et al., 2016).

To filter the Pareto front solutions based on objective performance, following the proposed framework, performance was compared to the two baseline models presented in Section 2.3.4.1. The two horizontal, dashed black lines show the benchmark metrics values calculated for average absolute clumpiness error, where the line of  $\Delta\text{CLU} = 0.060$  corresponding to the random constraint match benchmark metric value, and the line of  $\Delta\text{CLU} = 0.018$  corresponding to the growing clusters benchmark metric value. Of the parameter sets obtained, 74 resulted in superior performance to the random constraint match neutral model, and 50 resulted in superior

performance to the growing clusters neutral model. All solutions were considered valid with respect to FKS, as all had FKS agreement values greater than 0.


Figure 2.6 also shows the importance of considering model stochasticity as part of the automated adjustment process. The independent simulation objective values vary considerably from the average performance value, shown by the black line, with a noticeable increase in the spread of the independent solutions from the average for Pareto front solutions with higher FKS values. This variation is attributed to the influence of the stochastic factor which characterises the degree of randomness that is built into the land-use allocation for each model realisation (see transition potential definition in Equation 2.3) This highlights the need to consider model stochasticity appropriately when using automatic parameter tuning methods, as not doing so would result in over-calibration to a single model realisation.

#### 2.4.2 *Simulated output maps*


Following the assessment of the obtained Pareto front, this section presents an evaluation of a subset of the Pareto front solutions obtained following benchmark testing. The simulated output maps generated by certain parameter sets corresponding to the selected Pareto front solutions, shown by the black dots in Figure 2.6, were assessed to obtain a better understanding of the nature of the trade-off between the two objectives, by considering how the trade-off is reflected in the simulated output maps. Five Pareto front solutions were evaluated, with metric values given in Table 2.1. The selected solutions correspond to parameter sets that resulted in the lowest average values of FKS agreement, and lowest average values of  $\Delta\text{CLU}$ , the highest average value of FKS agreement with the highest average values of  $\Delta\text{CLU}$  that passed the benchmark test, and three intermediate solutions, approximately equally spaced with respect to the map comparison metrics. The corresponding simulated output maps are presented in Figure 2.7. For each figure, an enlarged version of the area in the black square in the left-hand panel is shown in the right-hand panel to highlight an area of significant change.

*Table 2.1. Metric values of Pareto front solutions for which simulated output maps were evaluated*

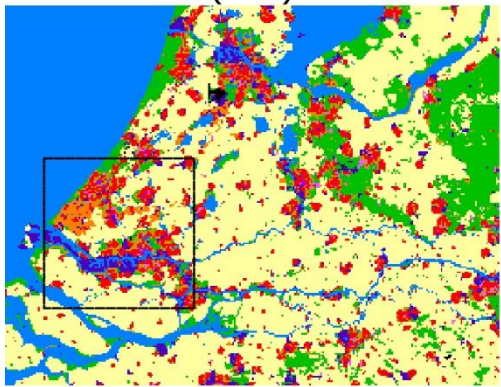
<b>Pareto Front solution</b>	<b>FKS</b>	<b><math>\Delta\text{CLU}</math></b>	<b>Corresponding map</b>
1	0.173	0.002	Figure 2.7b
9	0.181	0.003	Figure 2.7c
23	0.190	0.006	Figure 2.7d
36	0.201	0.009	Figure 2.7e
50	0.209	0.017	Figure 2.7f




(a-1)



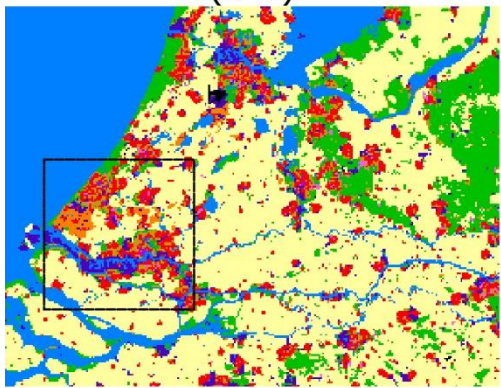
(a-2)




(b-1)




(b-2)




(c-1)



(c-2)



(d-1)



(d-2)



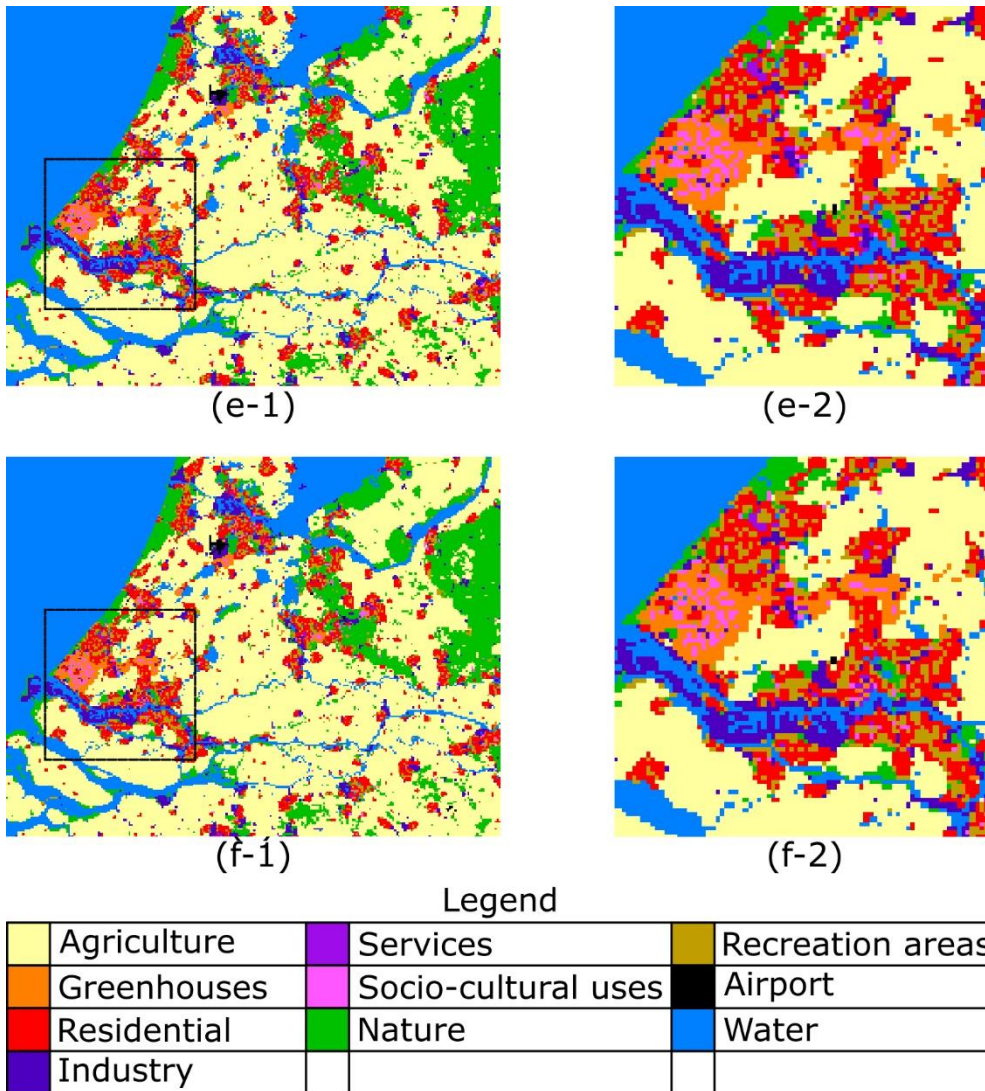


Figure 2.7. Comparison of selected simulated output maps corresponding to Pareto front solutions listed in Table 2.1 with data (a-1 and a-2). The maps are organised by increasing FKS values.

The variation between the simulated output maps can be largely attributed to the amount of interspersion observed between the different classes. The simulated output maps for parameter sets that achieved superior  $\Delta\text{CLU}$  performance, shown in Figures 2.7b and 2.7c, exhibit larger clusters of the socio-economic land-use classes, such as *residential* (red), *industrial* (purple), and *recreation areas* (dark-yellow). The simulated output maps for parameter sets that achieved superior FKS performance, such as Figures 2.7d, 2.7e and 2.7f, exhibited an increasingly large amount of interspersion of these land-use classes. The most notable examples can be seen with the interspersion of *recreation areas* amongst clusters of *residential* land-use throughout the map, also shown in Figures 2.7d-2, 2.7e-2 and 2.7f-2. There is also a growing presence of *socio-cultural uses* (pink) interspersed amongst *greenhouses* (orange) for



these simulated output maps, and a large cluster comprised of *greenhouses* and *residential* land uses that grows with improving FKS performance.

From visual inspection of the simulated output maps and comparison with the data, certain solutions could be rejected, despite passing the benchmark test, due to the patterns generated being too inconsistent with the data. For example, the solutions corresponding to maps shown in Figure 2.7d, 2.7e and 2.7f, are considered unsuitable due to the over-interspersion of the socio-economic land-use classes, which are not observed in the data. It appears that the use of FKS placed an over-emphasis on generating transitions, which resulted in the patterns observed in the simulated output maps being inconsistent with the data.

To evaluate this, the Kappa and Fuzzy Kappa metrics, more commonly used locational agreement metrics in calibration studies (e.g. Cao et al. (2015), Chaudhuri and Clarke (2013a), García et al. (2013), Wickramasuriya et al. (2009), Hagen-Zanker (2009)), that consider the entire landscape, were calculated. Values between 0.758 and 0.789 for Kappa and 0.828 and 0.854 for Fuzzy Kappa were obtained for the solutions generated. For Kappa, a benchmark value of 0.789 was calculated for the growing clusters neutral model, which is outperformed by only one solution, and 0.774 for the random constraint match neutral model, which is outperformed by 19 solutions, which corresponded to solutions that generated lower FKS values. For Fuzzy Kappa, a benchmark value of 0.869 was calculated for the growing clusters neutral model, which is not outperformed by any of the solutions, and 0.828 for the random constraint match neutral model, which is outperformed by 71 solutions. Such performance, with fewer solutions outperforming baseline metric values, is consistent with the over-emphasis on the agreement of transition cells with the use of the FKS metric to measure locational agreement.

#### **2.4.3 Parameter variation along Pareto front**

This section presents an analysis of the parameters corresponding to the different Pareto front solutions (coloured by solution as shown in Figure 2.6), to evaluate the interplay between different parameters and determine if this is consistent with process understanding, and to provide insight into how the preference for the different objectives was achieved through different parameter value settings. Given the large number of parameters in the LUCA model that were tuned (238), key parameters corresponding to the Pareto front solutions evaluated in Section 2.4.2 have been selected for detailed consideration: The inertia parameters for each land-use class; the conversion parameters for transitions to the class *residential*; the tails of the

interaction rules for transitions to the class *residential*; the interactions (conversion parameters and tails) between *greenhouses* and *socio-cultural uses*; and the distance decay parameters for the accessibility.

Figure 2.8 shows the inertia parameter values (the  $c$  value in Equation 2.6) for the self-influence rules, ordered and coloured by the Pareto front solution. As can be seen, there are trends in the parameters generated by the optimisation process. For the land-use classes *services* (SER) and *socio-cultural uses* (SCU), the optimisation process resulted in high influence values, relative to the values for the other land-use classes, across all five solutions evaluated. Hence, it can be concluded that both objectives were optimised when these land-use classes were relatively inert. For the other classes, the optimisation process generated a *residential* (RES) inertia value that was relatively high for solutions that favoured  $\Delta$ CLU (i.e. solutions 1 and 9), whereas the optimisation process resulted in all other classes having a relatively low inertia value, that was relatively constant between classes. However, for solutions that balance both objectives (solution 23), the inertia for *residential* drops to a relatively low value whilst the classes *industry* (IND), *nature* (NAT) and *recreation areas* (REC) jump to relatively higher values, a trend that persists for solutions favouring FKS (solutions 36 and 50). This suggests that, to generate better FKS values, there must be more transition from the class *residential*. This is reflected in data and confusion matrices for each solution (Supplementary material 2A). However, such a result is in conflict with common understanding of land-use dynamics, where the large initial investment required for urban classes, such as residential (Van Vliet et al., 2013b), means socio-economic land-use classes generally have relatively high inertia values, given the high associated transition costs. Additionally, the high inertia of the class *nature* also conflicts with common understanding of land-use dynamics, as natural areas are more often preserved due to zoning regulations (e.g. Van Delden et al., 2007) as opposed to being inert, as shown in Figure 2.8. Hence, the Pareto front solutions analysed exhibit certain inertia values that are consistent with expectation, such as *residential* for solution 1 and *industry* for solution 50. Other inertia values, such as the *industry* for solution 1 as well as *residential* and *nature* for solution 50, are not.




Figure 2.8. Variation of inertia parameter values for self-influence neighbourhood rules across Pareto front solutions

Figure 2.9 shows the values of the conversion parameter (the  $c$  value in Equation 2.6) for the interaction rules for different land-use classes converting to the class *residential*, ordered and coloured by the Pareto front solution. Generally, solutions favouring FKS exhibited higher conversion values, to generate more conversions to the class *residential*, which were exhibited in the data (see Supplementary material 2A). The trends shown in Figure 2.9 were also found for other conversion parameters. As shown, the optimisation process resulted in certain conversion parameters that did not exhibit any particular trend across the solutions, for example, the conversion from *industry* to *residential*, which has a relatively constant value for solutions 1 and 9 that favour  $\Delta$ CLU and solutions 36 and 50 that favour FKS, but a negative value for solution 23. A trend that was observed was that the optimisation process generated parameters that changed sign from negative to positive, as Pareto front solutions shifted from favouring  $\Delta$ CLU to FKS, to promote certain conversions. Conversion from *greenhouses*, *recreation areas*, and *services* to *residential* are negative for solutions 1 and 9 that favour  $\Delta$ CLU, preserving the landscape pattern by minimising potential conversions. Moving to solutions 23, 36 and 50, the optimisation process has resulted in positive parameter values, with increasing magnitude, to promote more conversions.




Figure 2.9. Variation of the conversion parameter influence values across Pareto front solutions for transitions from all land-use classes to the class residential

The tail parameters (the calculated  $y(x)$  value in Equation 2.6) corresponding to the neighbourhood rules exhibited similar trends to the conversion parameters, examples of which are illustrated for the influence of different land-use classes in the neighbourhood of the class *residential*, presented in Supplementary material 2B. There were certain neighbourhood rules that were tuned to relatively constant forms across the solutions, such as the weak attractive influence for the presence of *industry* in *residential* locations, likely due to both being socio-economic land-use classes. Other rules exhibited a trend across the solutions of an increasing attractive or repulsive strength, an example of which is shown in Figure 2.10, for the attraction of *residential* location in the neighbourhood of other *residential* locations. As shown, the attractive influences increase from relatively low values for solutions favouring  $\Delta$ CLU to relatively high values for solutions favouring FKS. There were also rules, such as the presence of *nature* in the neighbourhood of *residential* land-uses, which displayed no trend across the solutions. Solutions 1 and 9 that favoured  $\Delta$ CLU, and solutions 36 and 50 that favoured FKS, both featured a repulsive influence, but solution 23, which balances the two objectives, features a weak attractive influence.




Figure 2.10. Variation of the neighbourhood rule tails across Pareto front solutions for the influence of the class residential in the neighbourhood of the class residential

Examination of neighbourhood rules also provides insight into the unexpected model behaviour exhibited in the results maps, such as the large clusters of *socio-cultural uses* and *greenhouses* interspersed between each other in solutions 23, 36 and 50, highlighted in Figures 2.7d, 2.7e and 2.7f. Such behaviour could either result from higher conversion values, the  $c$  parameter, or a strong, attractive influence between the two classes, as a result of the  $a$  and  $b$  parameters (Equation 2.6). The conversion values, as well as the influence values at a distance of one, are presented Table 2.2, to indicate the resultant influence different land-use classes have on each other across the different solutions. As shown, the conversion values for transition from *greenhouses* to *socio-cultural uses* are all positive, facilitating such transitions across all solutions. However, only solutions 23, 36 and 50 had positive conversion values for transition from *socio-cultural uses* to *greenhouses*. A more likely explanation of the resultant behaviour is provided by the neighbourhood rule tails, indicated by the influence values at a distance of one across the five solutions. Solutions that preference  $\Delta$ CLU have a repulsive influence (i.e. a negative influence value) between both land-use classes, whilst solutions that preference FKS have an attractive influence between both land-use classes, producing a mutual attraction that explains the observed simulated output maps.

Table 2.2. Variation of conversion and influence at distance 1 of Pareto front solutions for conversion from greenhouses to socio-cultural uses and from socio-cultural uses to greenhouses

Rule	Solution	Conversion value	Influence at distance 1
From greenhouses to socio-cultural uses	1	80.07	-6.37
	9	58.04	-35.44
	23	93.88	201.37
	36	73.90	207.73
	50	91.03	210.88
From socio-cultural uses to greenhouses	1	-34.88	-366.20
	9	-28.37	-537.10
	23	36.40	114.23
	36	55.99	146.00
	50	81.46	168.97

To illustrate this further, the  $a$  parameters for the conversion from *greenhouses* to *socio-cultural uses*, which control the type (attractive or repulsive) and partially the strength of the influence that the presence of the class *greenhouses* has in the neighbourhood of cells transitioning to *socio-cultural uses*, is shown in Figure 2.11 for all Pareto front solutions, ordered by increasing FKS value. As shown, there is a clear point, at approximately solution 20, where the  $a$  parameter changes from a negative to a positive value, changing the relationship from repulsive to attractive. Given the trend shown in Figure 2.11, this appears to be a driver in generating solutions that preference FKS agreement. Similar behaviour was also found for the  $a$  parameter controlling the influence the class *socio-cultural uses* has in the neighbourhood of cells transitioning to *greenhouses*.

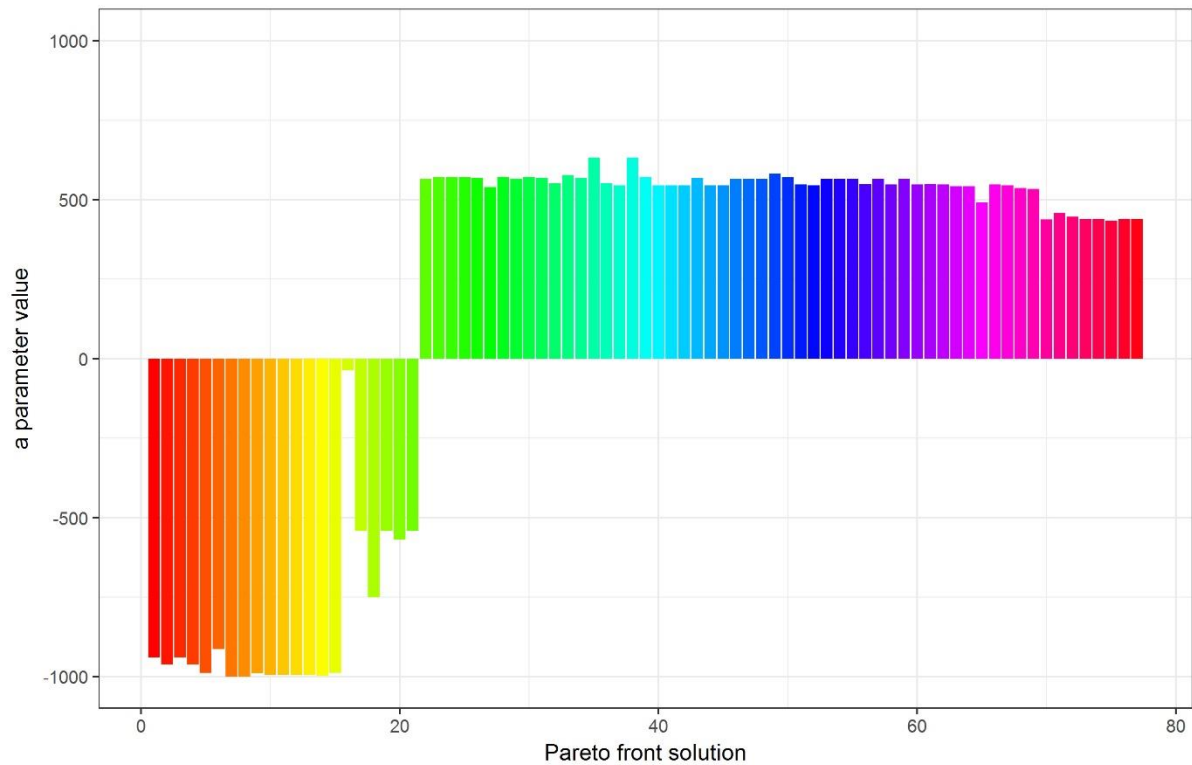


Figure 2.11. Variation of the *a* parameter value across the Pareto front solutions, ordered by increasing FKS, for the influence of greenhouses in the neighbourhood of socio-cultural uses

Figures 2.12 and 2.13 show the distance-decay parameter values for motorways and other roads, respectively, ordered and coloured by the selected solution. As shown, the optimisation process resulted in several trends across the Pareto front solutions. Some parameters were tuned to relatively constant values across the solutions, such as *industry* and *services* for motorways, and *recreation areas* and *socio-cultural uses* for other roads. Parameters for other classes were tuned to either relatively high or relatively low values, depending on the relative trade-off between the two objectives. For example, the distance decay for *greenhouses* and *residential* increased from a relatively low value for solutions favouring  $\Delta$ CLU to a relatively high value for solutions that favour FKS for both motorways and other roads, which is reflected in the simulation maps with the increasingly growing clusters of inter-connected patches of these land-uses between the two concentrated urban areas. In Figure 2.7b-2, for solutions with a preference for  $\Delta$ CLU, there are two distinct clusters of urban land-uses, most notably *residential*, that are sparsely connected, whereas in Figure 2.7f-2, for solutions with a preference for FKS, the two patches of residential area are essentially connected by the land-use classes *residential* and *greenhouses*, in an area where both motorways and other roads are present.




Figure 2.12. Accessibility distance-decay parameter values across Pareto front solutions for motorways




Figure 2.13. Accessibility distance-decay parameter values across Pareto front solutions for other roads



#### **2.4.4 Final parameter set**

Following the evaluation of the identified Pareto optimal solutions, considering the simulated output maps, parameters, and the influence of the objectives, the final parameter set recommended corresponds to Pareto front solution 1, which generated the solution map shown in Figure 2.7b (see Supplementary material 2C for parameters). This Pareto front solution outperformed the largest number of benchmark tests, and generated plausible solution maps. The parameters obtained more often aligned with general understanding of land-use change processes, although more neighbourhood rules were included than would be for a conventional manual calibration, though this was true for all solutions generated. Solutions that correspond to higher values of FKS agreement, as discussed, were not recommended because, based on the output maps and parameters, they did not result in realistic performance.

### **2.5 Discussion**

This section considers the results obtained from the optimisation process, with additional interpretation considering how the structure of the framework influences the results obtained.

#### **2.5.1 Enhanced understanding of applied metrics**

As the metrics drive the optimisation process, it is important to understand what impact different metrics have on the parameters and land-use maps obtained. With regard to the clumpiness metric, this was aggregated from class level into a single objective for optimisation purposes, which was achieved by taking the average. Consequently, it is important to understand the impact this averaging has on the results obtained. This was achieved by analysing the error for each class for the five solutions considered, as shown in Figure 2.14, with the colours corresponding to the different classes, and the average shown by the black bar on the left-hand side (for each solution). As shown in the figure, solutions 1 and 9 have relatively equal error across each class, which is captured well by the average. However, moving to solutions 23, 36 and 50, it is observed that there is a large variation in the errors across the classes, especially in solutions 23 and 36, where one class (*services* and *nature*, respectively) has a much higher error relative to the other classes.

These results suggest potential improvements could be achieved by using a different aggregation strategy for converting the class level clumpiness errors into a single value, for example, by taking the maximum, as this could potentially penalise the outputs more appropriately to generate patterns more consistent with the observed data. As the approach uses multi-objective optimisation, the individual class-level errors could also be considered as

independent objectives, which for this case-study would require eight objectives, including locational agreement and the seven clumpiness error values, generating an 8-dimensional objective space.




Figure 2.14. Class level absolute clumpiness error variation across the Pareto front for five selected solutions

With regard to the FKS measure, its use as a measure of locational agreement may have resulted in an over-emphasis on the agreement of transitions between the simulated output and the data (i.e. the corresponding actual land-use map), which is potentially not appropriately balanced with the agreement of inertia. This was confirmed by assessing the solutions with the Kappa and Fuzzy Kappa statistic (see Table 2.3). As shown in this table, the values with lower FKS correspond to higher Kappa and Fuzzy Kappa values. As previously mentioned, all solutions outperform the FKS benchmark (by being greater than 0), but only solution 1 outperforms the benchmark Kappa value, and no solutions outperform the benchmark Fuzzy Kappa value, generated by a growing clusters neutral model. Solutions 1 and 9 outperform the benchmark Kappa value generated by using a random constraint match neutral model, and all solutions outperform the benchmark Fuzzy Kappa values generated by using a random constraint match neutral model.

Table 2.3. A comparison of FKS, Kappa, and Fuzzy Kappa values across five selected Pareto front solutions

<b>Solution no.</b>	<b>FKS</b>	<b>Kappa</b>	<b>Fuzzy Kappa</b>
1	0.173	0.789	0.854
9	0.181	0.784	0.852
23	0.190	0.765	0.835
36	0.201	0.761	0.832
50	0.209	0.758	0.829
GC Neutral model	-	0.789	0.869
RCM Neutral model	-	0.773	0.828

Unlike for the landscape pattern structure metric, where improvement could be achieved with alteration of the aggregation method, the issues presented with the use of the FKS statistic are inherent to the metric, and cannot be adjusted without modifying the underlying calculation of the metric. The Fuzzy Kappa or Kappa statistic could potentially be used as an alternative, but this would likely introduce the opposite problem of the optimisation being too heavily weighted towards inertia within the measure of locational agreement. One solution may be to use both metrics as part of multi-objective optimisation, which would give more weight to the inertia (a land-use transition to the same land-use class) used in the calculation of FKS. Further investigation of the metrics used is a clear avenue for further research to improve the outputs of automated parameter tuning methods based on multi-objective optimisation.

### **2.5.2 *Balancing automatic and heuristic analysis in calibration***

The application of the multi-objective optimisation framework highlights the need to balance automatic and heuristic analysis. Too often, metrics quantifying different aspects of the fit between simulated output and data are used exclusively to evaluate the quality of a LUCA model calibration method. As covered in the results (Sections 2.4.2 and 2.4.3), evaluation can be improved with the incorporation of heuristic analysis, incorporating domain knowledge to evaluate the output solution maps and parameters.

The heuristic analysis also provides insights into the potential improvement of the calibration framework with the incorporation of additional process knowledge. Though many objectively valid solutions are produced, certain parameters are not necessarily consistent with those expected from process understanding. The incorporation of additional process knowledge regarding certain parameters presents a clear avenue to improving the optimisation procedure, by reducing the search-space and computational effort required, resulting in faster solution convergence, and parameters that are more consistent with expectations.

In turn, using an automatic approach also provides insights into the model structure and its behaviour. Certain model parameter values, though inconsistent with heuristic expectation, still generated results that were acceptable according to the metrics used. This provided insight into what aspects of the parameterisation and model results should be captured by the objectives (Section 2.5.1) and also showed the importance of providing plausible boundaries for model parameters. Further work on the automatic calibration, including a thorough analysis of the solutions obtained, presents a path for future research to improve the understanding of LUCA model calibration, the model itself and the interplay of factors driving land use change.

## **2.6 Conclusions and recommendations**

Accurate and effective calibration of LUCA models is essential to allow for application to the study region of interest. This paper presents an automatic calibration framework specifically designed for such LUCA models. The framework is generic, allowing for substitution of the various components (e.g. the LUCA model, metrics, and the optimisation algorithm), utilises multi-objective optimisation to allow for the exploration of trade-offs between the selected objectives, appropriately considers the inherent stochasticity included in LUCA models, and facilitates increased computational efficiency through its implementation.

The capability of the generic framework has been illustrated with an application to the Randstad region of the Netherlands. The case study was implemented with parallel computing, to reduce the parameter adjustment computing time from the order of hundreds of days to the order of days. The approach generated 77 possible model parameter sets that resulted in Pareto optimal solutions based on the selected objective metrics, of which 50 were considered plausible. Following evaluation of both the simulated output and parameters against process knowledge, a final parameter set was recommended.

The framework provided crucial insights into the influence the metrics had on driving the optimisation process, and how more realistic performance could potentially be achieved with the alteration to the optimisation objectives. The results are also driven by the genetic algorithm parameters used. As with other applications of genetic algorithms to the automatic parameter adjustment stage of calibration of LUCA models (e.g. Jafarnejhad et al., 2016) it would be meaningful to investigate the influence these parameters had on the output obtained, though this would depend on computational resource availability.

The framework highlights the importance of heuristic interpretation when evaluating the result map and obtained parameters, and the potential of such methods to enhance LUCA model

understanding. The demonstrated application of the framework and its generic nature show promising potential for future applications of LUCA models to support long term planning and policy development.

## **2.7 Acknowledgements**

The authors wish to acknowledge the financial support from the Bushfire and Natural Hazards Cooperative Research Centre, made available by the Commonwealth of Australia through the Cooperative Research Centre program. The authors also wish to acknowledge that this work was supported with supercomputing resources provided by the Phoenix High Performance Computing Services at the University of Adelaide.

## 2.8 Supplementary material

### 2.8.1 Supplementary material 2A: Confusion matrices

This section contains confusion matrices for the data and the six solution maps presented in Section 2.4.2, showing a count of land-use class presence per cell between 1989 and 2000. These tables highlight the volume and type of land-use transitions across the different Pareto front solutions.

Table 2.A1. Confusion matrix for data between 1989 and 2000

		Map 2										
	LUC	AGR	GRE	RES	IND	SER	SCU	NAT	REC	AIR	WAT	Total
Map 1	AGR	26128	152	559	453	12	17	641	336	12	188	28498
	GRE	81	349	25	9	1	2	1	1	0	5	474
	RES	160	3	3511	35	18	37	34	123	0	22	3943
	IND	67	1	45	1205	7	10	43	28	7	65	1478
	SER	3	0	32	37	68	12	6	9	0	4	171
	SCU	13	0	55	17	2	212	39	20	0	3	361
	NAT	92	1	19	26	0	3	5814	129	0	50	6134
	REC	105	6	81	28	2	21	141	1229	0	44	1657
	AIR	34	0	0	0	0	0	3	0	31	0	68
	WAT	110	2	23	85	0	3	56	37	0	1685	2001
	Total	26793	514	4350	1895	110	317	6778	1912	50	2066	44785

Table 2.A2. Confusion matrix for Pareto optimal solution with  $FKS=0.173$  and  $\Delta CLU = 0.002$  between 1989 and 2000

		Map 2										
	LUC	AGR	GRE	RES	IND	SER	SCU	NAT	REC	AIR	WAT	Total
Map 1	AGR	26529	41	426	447	0	0	787	268	0	0	28498
	GRE	0	473	0	0	0	0	0	1	0	0	474
	RES	0	0	3921	0	0	0	0	22	0	0	3943
	IND	30	0	1	1426	0	0	1	20	0	0	1478
	SER	11	0	1	1	110	0	2	46	0	0	171
	SCU	25	0	1	5	0	317	0	13	0	0	361
	NAT	64	0	0	15	0	0	5962	93	0	0	6134
	REC	181	0	0	1	0	0	26	1449	0	0	1657
	AIR	0	0	0	0	0	0	0	0	68	0	68
	WAT	0	0	0	0	0	0	0	0	0	2001	2001
	Total	26840	514	4350	1895	110	317	6778	1912	68	2001	44785

Table 2.A3. Confusion matrix for Pareto optimal solution with  $FKS=0.181$  and  $\Delta CLU = 0.003$  between 1989 and 2000

		Map 2											
		LUC	AGR	GRE	RES	IND	SER	SCU	NAT	REC	AIR	WAT	Total
Map 1	LUC												
	AGR	26499	108	475	398	0	0	807	211	0	0		28498
	GRE	14	406	4	27	0	0	6	17	0	0		474
	RES	0	0	3870	0	0	0	0	73	0	0		3943
	IND	0	0	0	1448	0	0	1	29	0	0		1478
	SER	15	0	1	4	110	0	2	39	0	0		171
	SCU	26	0	0	3	0	317	2	13	0	0		361
	NAT	117	0	0	15	0	0	5913	89	0	0		6134
	REC	169	0	0	0	0	0	47	1441	0	0		1657
	AIR	0	0	0	0	0	0	0	0	68	0		68
	WAT	0	0	0	0	0	0	0	0	0	2001		2001
	Total	26840	514	4350	1895	110	317	6778	1912	68	2001		44785

Table 2.A4. Confusion matrix for Pareto optimal solution with  $FKS=0.190$  and  $\Delta CLU = 0.006$  between 1989 and 2000

		Map 2											
		LUC	AGR	GRE	RES	IND	SER	SCU	NAT	REC	AIR	WAT	Total
Map 1	LUC												
	AGR	26182	150	770	381	1	13	666	335	0	0		28498
	GRE	26	363	17	3	0	50	1	14	0	0		474
	RES	252	1	3477	6	4	3	5	195	0	0		3943
	IND	0	0	1	1466	1	0	0	10	0	0		1478
	SER	25	0	21	8	101	0	3	13	0	0		171
	SCU	56	0	36	5	0	250	1	13	0	0		361
	NAT	21	0	19	13	1	0	6025	55	0	0		6134
	REC	278	0	9	13	2	1	77	1277	0	0		1657
	AIR	0	0	0	0	0	0	0	0	68	0		68
	WAT	0	0	0	0	0	0	0	0	0	2001		2001
	Total	26840	514	4350	1895	110	317	6778	1912	68	2001		44785

Table 2.A5. Confusion matrix for Pareto optimal solution with  $FKS=0.201$  and  $\Delta CLU = 0.009$  between 1989 and 2000

		Map 2											
		LUC	AGR	GRE	RES	IND	SER	SCU	NAT	REC	AIR	WAT	Total
Map 1	LUC												
	AGR	26145	183	946	358	0	24	634	208	0	0		28498
	GRE	40	323	25	6	0	73	0	7	0	0		474
	RES	275	5	3257	6	0	5	10	385	0	0		3943
	IND	0	0	3	1466	0	1	0	8	0	0		1478
	SER	9	1	24	10	110	0	2	15	0	0		171
	SCU	61	0	68	6	0	212	3	11	0	0		361
	NAT	0	0	4	1	0	0	6077	52	0	0		6134
	REC	310	2	23	42	0	2	52	1226	0	0		1657
	AIR	0	0	0	0	0	0	0	0	68	0		68
	WAT	0	0	0	0	0	0	0	0	0	2001		2001
	Total	26840	514	4350	1895	110	317	6778	1912	68	2001		44785

Table 2.A6. Confusion matrix for Pareto optimal solution with  $FKS=0.210$  and  $\Delta CLU = 0.019$  between 1989 and 2000

		Map 2											
		LUC	AGR	GRE	RES	IND	SER	SCU	NAT	REC	AIR	WAT	Total
Map 1	LUC												
	AGR	26031	203	938	365	0	22	650	289	0	0	0	28498
	GRE	69	289	31	7	0	64	2	12	0	0	0	474
	RES	357	13	3250	12	0	4	6	301	0	0	0	3943
	IND	0	0	6	1453	0	1	0	18	0	0	0	1478
	SER	11	4	22	6	110	0	2	16	0	0	0	171
	SCU	59	1	46	12	0	224	3	16	0	0	0	361
	NAT	0	0	14	0	0	0	6044	76	0	0	0	6134
	REC	313	4	43	40	0	2	71	1184	0	0	0	1657
	AIR	0	0	0	0	0	0	0	0	68	0	0	68
	WAT	0	0	0	0	0	0	0	0	0	2001	0	2001
	Total	26840	514	4350	1895	110	317	6778	1912	68	2001	0	44785



### 2.8.2 Supplementary material 2B: Neighbourhood tail influence figures

This section contains a set of figures corresponding to the different neighbourhood rules describing the influence different land-use classes exert on the class *residential* for the solutions analysed in detail.




Figure 2.B1. Neighbourhood rule tail for conversion from agriculture to residential




Figure 2.B2. Neighbourhood rule tail for conversions from greenhouses to residential




Figure 2.B3. Neighbourhood rule tail for conversions from industrial to residential




Figure 2.B4. Neighbourhood rule tail for conversions from nature to residential




Figure 2.B5. Neighbourhood rule tail for conversions from recreation areas to residential




Figure 2.B6. Neighbourhood rule tail for conversions from residential to residential




Figure 2.B7. Neighbourhood rule tail for conversions from socio-cultural uses to residential




Figure 2.B8. Neighbourhood rule tail for conversions from services to residential

### 2.8.3 Supplementary material 2C: Final model parameterisation

Table 2.C1. Inertia/Conversion parameters

From/To	GRE	RES	IND	SER	SCU	NAT	REC
AGR	94.62	81.10	35.30	48.96	-26.93	20.45	15.88
GRE	2608.65	-40.30	80.86	-35.52	58.04	-38.79	81.79
RES	58.02	4142.51	-29.37	69.19	45.49	79.16	71.61
IND	2.56	45.18	2504.04	-37.57	31.13	-36.41	-46.91
SER	25.18	-31.74	-43.44	4508.40	71.34	4.68	48.02
SCU	-28.37	29.59	-4.36	-7.58	4124.71	33.40	37.45
NAT	72.61	81.08	37.25	70.26	-38.65	2360.85	100.00
REC	69.41	-32.05	49.51	-46.85	-3.60	89.57	2957.07
AIR	53.73	88.52	17.15	57.37	-18.60	53.27	69.24
WAT	-49.51	31.98	70.93	44.79	88.21	62.55	-31.23

Table 2.C2. Neighbourhood rule a parameters

From/To	GRE	RES	IND	SER	SCU	NAT	REC
AGR	-179.73	-704.44	-720.20	-374.93	86.68	-657.23	-680.28
GRE	707.69	247.79	-466.96	233.29	-990.06	-237.31	-811.14
RES	-996.46	369.06	105.07	110.30	-485.26	-815.79	817.88
IND	127.06	184.53	137.92	234.48	768.52	487.25	-56.00
SER	-678.64	897.13	47.45	-845.38	-447.53	69.96	117.35
SCU	-565.31	394.57	985.62	-83.81	171.47	-336.22	778.66
NAT	195.83	-875.47	-246.59	-844.79	683.76	171.47	-988.66
REC	-283.93	85.62	-663.00	-302.89	-691.96	602.16	613.24
AIR	-12.20	-942.51	851.11	-461.78	134.85	-984.62	-50.08
WAT	-63.18	-636.50	-433.31	-266.81	-898.21	776.96	119.15

Table 2.C3. Neighbourhood rule b parameters

From/To	GRE	RES	IND	SER	SCU	NAT	REC
AGR	4.32	3.58	1.37	0.66	4.06	4.91	0.96
GRE	3.21	4.47	4.62	2.02	3.33	2.32	3.58
RES	3.58	1.87	4.81	4.73	4.91	3.23	1.21
IND	3.29	2.53	0.73	3.93	4.13	1.99	4.48
SER	1.54	2.32	4.10	3.30	4.03	2.41	0.19
SCU	0.05	3.32	4.73	2.82	3.16	3.31	3.87
NAT	2.89	1.96	1.16	3.90	4.70	3.16	2.43
REC	4.45	3.12	1.84	1.40	4.02	1.88	2.12
AIR	3.74	1.76	2.55	4.49	3.34	2.10	1.91
WAT	4.85	3.76	3.74	3.79	1.48	3.04	4.24

Table 2.C4. Accessibility parameters

<b>Land-use class</b>	<b>Motorway Distance Decay</b>	<b>Motorway weight</b>	<b>Other roads Distance Decay</b>	<b>Other roads weight</b>
GRE	14.33	0.48	0.01	0.46
RES	3.56	0.93	12.39	0.66
IND	7.00	0.38	13.42	0.59
SER	1.27	0.48	4.34	0.21
SCU	7.75	0.26	10.56	0.37
NAT	1.73	0.71	12.14	0.99
REC	8.40	0.19	3.20	0.13





### **3 Empirically derived method and software for semi-automatic calibration of Cellular Automata land-use models**

Newland, C.P., Zecchin, A.C., Maier, H.R., Newman, J.P. & Van Delden, H. (Submitted). Empirically derived method and software for semi-automatic calibration of Cellular Automata land-use models. *Environmental Modelling & Software* (Adapted for consistency).

## Statement of Authorship

Title of Paper	Empirically derived method and software for semi-automatic calibration of Cellular Automata land-use models		
Publication Status	<input type="checkbox"/> Published	<input type="checkbox"/> Accepted for publication	
	<input checked="" type="checkbox"/> Submitted for publication	<input type="checkbox"/> Unpublished and Un-submitted work written in manuscript style	
Publication Details	Resubmitted following reviewer feedback to Environmental Modelling & Software		

## Principal Author

Name of Principal Author (Candidate)	Charles P. Newland		
Contribution to the Paper	Designed scope of study and experimental procedure, developed software, conducted experiments, performed analysis of results, wrote manuscript and acted as corresponding author.		
Overall percentage (%)	70%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.		
Signature		Date	28 <sup>th</sup> Feb 2018

## Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

- i. The candidate's stated contribution to the publication is accurate (as detailed above);
- ii. Permission is granted for the candidate to include the publication in the thesis; and
- iii. The sum of all co-author's contributions is equal to 100% less the candidate's stated contribution.

Name of Co-author	Aaron C. Zecchin		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript and response to reviewer comments.		
Signature		Date	28 <sup>th</sup> Feb 2018

Name of Co-author	Holger R. Maier		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript and response to reviewer comments.		
Signature		Date	28 <sup>th</sup> Feb 2018

Name of Co-author	Jeffrey P. Newman		
Contribution to the Paper	Assisted with software development. Reviewed manuscript.		
Signature		Date	28 <sup>th</sup> Feb 2018

Name of Co-author	Hedwig Van Delden		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript and response to reviewer comments.		
Signature		Date	28 <sup>th</sup> Feb 2018

# **Empirically derived method and software for semi-automatic calibration of Cellular Automata land-use models**

## **Abstract**

Land-use change models generally include neighbourhood rules to capture the spatial dynamics between different land-uses that drive land-use changes, introducing many parameters that require calibration. We present a process-specific semi-automatic method for calibrating neighbourhood rules that utilises discursive knowledge and empirical analysis to reduce the complexity of the calibration problem, and efficiently calibrates the remaining interactions with consideration of locational agreement and landscape pattern structure objectives. The approach and software for implementing it were tested on four case studies of major European cities with different physical characteristics and rates of urban growth. The implementation of the proposed approach explored the impact of a preference for different objectives, and outperformed benchmark models for both calibration and validation when a balanced objective preference was used. This research demonstrates the utility of process-specific calibration methods, and highlights how discursive knowledge can be integrated with automatic calibration to make the process more efficient.

### 3.1 Introduction

Land-use change models are used to understand the wide-ranging impacts of land-use changes, including the impact on the rate of greenhouse gas emissions (Li et al., 2017, Pogson et al., 2016), the balance of agricultural production with ecosystem preservation (Connor et al., 2015, Van Delden et al., 2010), and the influence of land-use policy on urban growth (Berberoğlu et al., 2016, Chaudhuri and Clarke, 2013a). Land-use changes are caused by many different, mutually influential bio-physical and socio-economic drivers (Lambin et al., 2001, Wang et al., 2011a) which must be captured effectively by land-use change models to generate realistic output. To model land-use changes effectively Cellular Automata (CA) were proposed to replicate land-use as a dynamic spatial system (Tobler, 1979). This required relaxing the conventional implementation of CA (Couclelis, 1985) to replicate fractal patterns of land-use changes consistent with urban evolution (White and Engelen, 1993a). Doing so led to the development of multiple Land Use Cellular Automata (LUCA) models, which have proliferated due to their simplicity, flexibility and intuitiveness (Santé et al., 2010).

An important aspect of modelling land-use changes is the consideration of spatial and temporal dynamics between different land-uses (Van Vliet et al., 2013b). Applications of CA to land-use modelling consider the composition of the neighbourhood. A strict CA only considers the neighbourhood composition of the geometrically closest set of cells for the consideration of spatial dynamics, implemented in numerous LUCA models (Li et al., 2013, Wu, 2002), including the popular generic modelling framework SLEUTH (Clarke et al., 1997, Silva and Clarke, 2002). However, other LUCA models, most notably those derived from White and Engelen (1993c) that use a transition potential as the mechanism for the allocation of land-use changes, use an extended neighbourhood (i.e. the set of cells within a certain cellular radius) for a more detailed consideration of spatial dynamics, which requires the use of neighbourhood rules that characterise the influence different land-use classes exert on each other relative to proximity.

While the use of neighbourhood rules allows for a more detailed replication of the spatial dynamics that exist between different land-use classes, it also increases the number of model parameters significantly, often by hundreds (Blecic et al., 2015, García et al., 2013). LUCA model application requires appropriate values of these parameters to be determined by means of calibration, which involves the initial setting of parameter values, the iterative adjustment of these values based on comparison of the model output with observations, and the selection of a final parameter set, for application to a specific case study for long term scenario analysis

(Newland et al., 2018a). The difficulty of the calibration processes is generally a function of the dimensionality of the parameter space, which is a function of the number and possible ranges parameters can take during the iterative adjustment process. Neighbourhood rules are the primary contributor to this high parameter dimensionality, and hence are often the main calibration parameters of these types of LUCA models (Engelen and White, 2008).

The conventional method of calibration is to manually adjust the neighbourhood rule parameters (García et al., 2013, Van Delden et al., 2012), which is a knowledge and time intensive process. Given the complexity of the calibration problem, the parameter dimensionality is often implicitly reduced when using such an approach based on the modeller's knowledge of the spatial dynamics driving land-use changes in a region (White et al., 1997), and the common forms of neighbourhood rules (Hagoort et al., 2008). Additionally, metrics can be used to objectively evaluate calibration performance. For LUCA models, there are two distinct types of metrics that measure different aspects of calibration performance (Newland et al., 2018a): locational agreement metrics, the match of pixels between simulated output and data, and landscape pattern structure metrics, the difference between simulated and observed patterns of land-use that infer the realism of land-use change processes. Objective measures of LUCA model performance have facilitated the development of automatic calibration methods, where parameters are iteratively adjusted automatically to improve LUCA model performance as quantified by the metrics used.

There are two general approaches to automatically calibrate neighbourhood rules. The first approach uses a population-based optimisation algorithm (Blecic et al., 2015, García et al., 2013, Newland et al., 2018a), where a population of solutions is generated (i.e. a number of different sets of neighbourhood rules), and adjusted based on some operators to improve the objective performance of the solutions over a number of iterations. Optimisation approaches are effective at generating multiple possible model parameterisations, but are computationally intensive, often requiring parallel computing resources for practical implementation (Blecic et al., 2015, Newland et al., 2018a). Such approaches also lack transparency, which can result in parameters that conflict with process understanding, despite producing a model with objectively good performance. The alternative to automatic calibration is the use of efficiency focussed methods targeted to the spatial dynamic processes in the transition potential model, aiming to generate a single set of calibrated neighbourhood rules that is consistent with discursive knowledge for a limited computational budget (Straatman et al., 2004, Maas et al., 2005, Van Vliet et al., 2013b), achievable using a desktop PC as opposed to a supercomputer.

Given the limited availability of supercomputing resources, such approaches, designated as process-specific, are valuable as a practical means of automatic calibration.

Despite a specific focus on neighbourhood rules, previous process-specific methods have not fully utilised discursive knowledge to generate neighbourhood rules consistent with such knowledge (Van Vliet et al., 2013b), and do not necessarily focus on the most important spatial interactions during calibration (Straatman et al., 2004). Previous methods have also only used a single metric of performance, not considering the implementation of multiple metrics to capture the two aspects of calibration performance previously discussed, and how these competing objectives impact on the resulting model performance.

Given the shortcomings of current process-specific calibration methods outlined above, while maintaining their relative computational efficiency and transparency compared with population-based optimisation algorithms, this research proposes a semi-automatic, process-specific calibration method. The aims of this paper are (i) to develop a calibration method that utilises process knowledge about meaningful interactions to facilitate efficient automatic calibration that allows for the consideration of both aspects of LUCA model performance for a limited computational budget, (ii) to investigate the utility of the proposed approach through application to several case studies, and (iii) to obtain insight into the choice of objectives and preference for them on the resultant calibrated model.

The remainder of this paper is organised as follows: To illustrate the complexity of the calibration problem, background information about neighbourhood rules is presented in Section 3.2. Section 3.3 presents the proposed approach for calibration, and Section 3.4 presents the implementation and computational testing regime used to evaluate it. The results are presented and discussed in Section 3.5, and the conclusions of this work are presented in Section 3.6.

### **3.2 Complexity of the calibration problem**

This section provides explicit details of the complexity of the calibration of neighbourhood rules, by defining the dimensionality of the neighbourhood rule parameter space for transition potential LUCA models. To do this, consider the transition potential land-use allocation mechanism, illustrated in Figure 3.1, which is described in detail below.

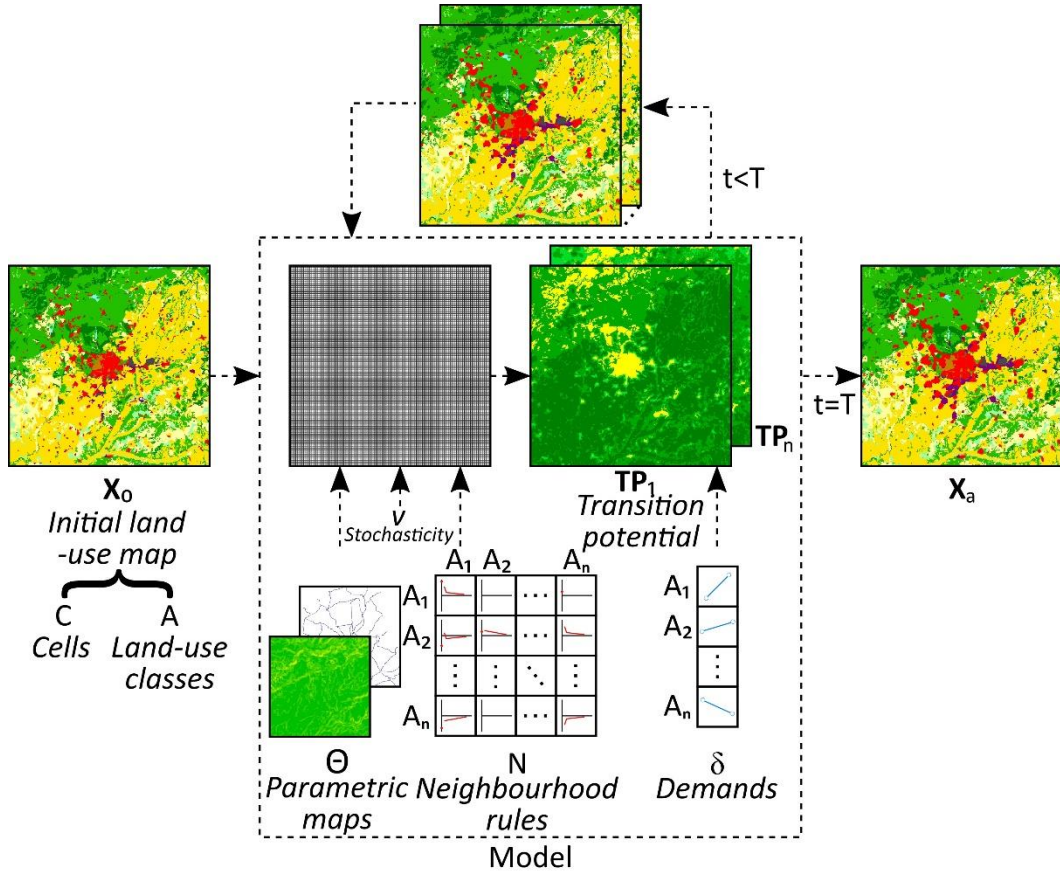


Figure. 3.1. Example of a transition potential allocation mechanism for a CA land-use model. Starting from an initial land-use map, the model allocates a transition potential to each cell (the gridded map) for each type of land-use based on a set of parametric maps, neighbourhood rules, and a stochastic component. The model then allocates a land-use class to each cell based on the potential and demand for each year ( $t$ ). This process is repeated until the final year of the simulation ( $t=T$ ), when the output map for that year is produced.

A transition potential based LUCA model can be defined as a dynamic process providing a sequence of land-use maps,  $X_1, X_2, \dots, X_T$ , given an initial land-use map  $X_0$ . Each land-use map is defined specifically as the set of land-use classes for each cell:

$$X_t = \{x_{c,t} \in A : c \in C\} \quad (3.1)$$

where  $x_{c,t}$  is the land-use class of cell  $c$  at time  $t$ ,  $A$  is the set of all land-use classes, and  $C$  is the set of all cells. For most applications, the set  $A$  can be partitioned as:

$$A = A_A \cup A_P \cup A_S \quad (3.2)$$

where  $A_A$  is the set of active land-use classes (i.e. the classes that are actively allocated by the model, generally urban classes such as residential and commercial land-use),  $A_P$  is the set of passive classes (i.e. classes that have low transition costs such as natural areas, and are allocated after the actively modelled classes) and  $A_S$  is the set of static classes (i.e. classes that occupy a



fixed location throughout the simulation, such as water bodies or airports). The set  $C$  can be partitioned following the same convention:

$$C = C_{A,t} \cup C_{P,t} \cup C_S \quad (3.3)$$

where  $C_{A,t}$  is the set of cells containing active land-use classes at time  $t$ ,  $C_{P,t}$  is the set of cells containing passive land-use classes at time  $t$ , and  $C_S$  is the set of cells containing static classes, which do not vary with time.

The evolution dynamics of the land-use is given by the functional relationship:

$$\mathbf{X}_{t+1} = f(\boldsymbol{\theta}_t, \mathbf{N}(\mathbf{X}_t), \mathbf{v}_t, \boldsymbol{\delta}) \quad (3.4)$$

where  $\boldsymbol{\theta}_t$  are parametric maps governing the effect of processes such as soil quality and zoning that influence land-use changes (which are independent of the land-use map),  $\mathbf{N}(\mathbf{X}_t)$  is the set of neighbourhood influence values that is dependent on the current state of the land-use map,  $\mathbf{v}_t$  is a stochastic perturbation term, included to capture the system uncertainty, and  $\boldsymbol{\delta}$  is the land-use demand (assuming the LUCA model is constrained) for each active land-use class. Each of  $\boldsymbol{\theta}_t$ ,  $\mathbf{N}(\mathbf{X}_t)$ , and  $\mathbf{v}_t$  are given for each cell in the land-use map, for each actively and passively modelled land-use class. For each time step  $t$  (i.e. each year) the functional relationship in Equation 3.4 describes how land-use is allocated to every cell in the map based on the transition potential, which is the potential for each specific cell to support each type of active or passive land-use class. For each active class, the cells with the highest transition potential are selected, until the demand is met, after which passive land-uses are allocated to the remaining cells (as these have no specified demand). A general form for determining transition potential is:

$$TP_{c,i,t} = \theta_{c,i,t} \cdot N_{c,i,t} \cdot v_t \quad (3.5)$$

where  $TP_{c,i,t}$  is the potential for cell  $c$  to support land-use  $i$  at time  $t$ ,  $N_{c,i,t}$  is the neighbourhood potential for cell  $c$  to support land-use  $i$  at time  $t$ , and is a function of the composition of the cellular neighbourhood of  $c$ , with each cell in the neighbourhood exerting some influence based on its class and distance from  $c$ . This has the general form:

$$N_{c,i,t} = \sum_{c' \in D(c)} w_{i,x_{c',t},d(c,c')} \quad (3.6)$$

where  $D(c)$  is the set of cells in the neighbourhood of cell  $c$ ,  $d(c,c')$  is the linear distance between the cells  $c$  and  $c'$ , and  $w_{i,j,d}$  is a neighbourhood weighting parameter that expresses the influence that a cell of land-use type  $j$  exerts on the potential for land-use class  $i$  at a linear

distance  $d$  between the two cells (note that in Equation 3.6, the subscripts of the general form of the weight term  $w_{i,j,d}$  are given by  $j$  equalling  $x_{c',t}$ , and  $d$  equalling  $d(c,c')$ ). Neighbourhood rules define the weight values, capturing distinct aspects of spatial dynamics based on the type of relationship being described and the distances being considered (Van Vliet et al., 2013b). This is illustrated in Figure 3.2, showing how weights are defined at different distances for different spatial dynamics. It is important to note that the weight values have no physical meaning, but derive meaning relative to each other.

In total, the set of neighbourhood interactions of class  $j$  on class  $i$  at distance  $d$  (symbolised by  $(i,j)_d$ ) are defined by the set:

$$P = \{(i,j)_d : i \in A_A, j \in A, d \in [0] \cup [1, d_{MAX}]\} \quad (3.7)$$

where  $P$  is the set of all interactions,  $A_A$  is the set of actively allocated land-use classes,  $A$  is the set of all land-use classes, and  $d_{MAX}$  is the maximum distance that defines the size of the neighbourhood. Interactions can be categorised as either self-influencing, when  $i$  and  $j$  are equal, (e.g. describing the influence of a residential class to attract more residential development), or across class interactions, when  $i$  and  $j$  are unequal (e.g. describing the influence of a residential class to attract industrial development). All interactions can also be described as either a point influence at distance zero, describing the influence of a cell on itself; or a neighbourhood tail influence for distances of one or greater, describing the remote influence of a class on a cell. Given these categorisations, there are four distinct types of spatial dynamics, examples of which are given in Figure 3.2, which are captured by neighbourhood interaction types (Van Vliet et al., 2013b), outlined below:

1. Inertia points,  $[(i,i)_0 \in P, i \in A_A]$ : These describe the persistence of a specific land-use class  $i$  to remain in its present location;
2. Conversion points,  $[(i,j)_0 \in P, i \in A_A, j \in A, j \neq i]$ : These describe the potential for a transition from one land-use class  $j$  to a different land-use class  $i$  at its present location;
3. Self-influence tails,  $[(i,i)_d \in P, i \in A_A, d \geq 1]$ : These describe the influence a land-use class  $i$  exerts on the same type of land-use  $i$  that is in the neighbourhood of a cell; and
4. Cross-influence tails,  $[(i,j)_d \in P, i \in A_A, j \in A, j \neq i, d \geq 1]$ : These describe the influence a land-use class  $j$  exerts on a different type of land-use  $i$  that is in the neighbourhood of a cell.




Figure 3.2. An example of how different neighbourhood rules and types of interactions are used to determine neighbourhood potential (in a single plane to the right of the cell marked with an  $x$ ), an inertia point and self-influence tail, which influences the potential for the cell of interest to remain residential (red), and a conversion point and cross-influence tail, which influences the potential for the cell to transition to industrial (purple)

The four distinct types of interaction mean the full set can be partitioned as:

$$P = P_{IP} \cup P_{CP} \cup P_{ST} \cup P_{CT} \quad (3.8)$$

where  $P_{IP}$  is the set of inertia point interactions,  $P_{CP}$  is the set of conversion point interactions,  $P_{ST}$  is the set of self-influence tail interactions, and  $P_{CT}$  is the set of cross-influence tail interactions. The total number of interactions considered is generally given by:

$$P = |A_A| \cdot |A| \quad (3.9)$$

Each tail interaction can be described by the weight at distance 1 multiplied by a function that describes the decaying neighbourhood tail interaction, that is:

$$w_{i,j,d} = w_{i,j,1} \cdot u(d) \text{ for } d \geq 1 \quad (3.10)$$

where  $u(d)$  is the functional form of the neighbourhood rule tail shape. The tail shape  $u(d)$  is a pre-specified function (e.g. linear, exponential) starting at  $u(1)$ , and decaying as  $d$  increases, resulting in a diminishing effect of a land-use class on another with increasing distance.

As each neighbourhood influence can feature four parameters, as shown in Figure 3.2 (e.g. a point weight at distance 0, and a piecewise linear relationship of tail weights for distance 1, a point of inflection, and a point where influence is set to zero), and as LUCA models can typically have 20 land-use classes where 10 are actively modelled, there are minimally 800 parameters that require calibration. Hence, the dimensionality of the neighbourhood rule parameter space is typically very high. It is for this reason that neighbourhood rules are the main focus of the calibration of LUCA models (Engelen and White, 2008), and why the calibration of neighbourhood rules is complex. The high dimensionality makes it difficult to

know which parameters to adjust to achieve an improved model calibration. This problem is further exacerbated by the limited calibration data available, as typically there exists only an initial data land-use map  $\hat{X}_0$  and one or two other data land-use maps  $\hat{X}_a$  and  $\hat{X}_b$  for some  $a$  and  $b \geq 1$ . The high parametric dimensionality also introduces potential issues with equifinality (Van Vliet et al., 2016), where the same calibration performance can be achieved by different sets of neighbourhood rules, despite the rules not necessarily being consistent with process knowledge.

### 3.3 Proposed approach

A conceptual outline of the proposed approach for the (semi) automatic calibration of neighbourhood rules is shown in Figure 3.3. The approach has been developed with two primary aims. The first is the reduction of the dimensionality of the neighbourhood rule parameter space, to mitigate the parametric dimensionality issues outlined previously, by identifying the key land-use interactions for consideration (symbolised by the set  $P'$ ). This is based on automating and formalising the process of interaction elimination that is common in manual calibration methods. The proposed approach is conducted in an objective, transparent and repeatable manner, however, it also allows for the manual elimination of interactions.

The second aim is the computationally efficient calibration of the remaining neighbourhood rule weighting parameters. An initial set of neighbourhood rule weighting parameters  $W'_{initial}$  is generated during the parameter categorisation stage, and is subsequently refined in the coarse adjustment stage to generate  $W'_{coarse}$ , a point determined with user input within the parameter space that results in objectively good performance that is consistent with process knowledge. Next, during the fine adjustment stage, the neighbourhood rule weighting parameters are individually tuned to generate a set of calibrated weighting parameters  $W'_{final}$ . Advantages of the proposed methods are that calibration is performed efficiently with a minimal number of model simulations (i.e. a computational budget achievable using a desktop PC as opposed to a supercomputer), the two key aspects of LUCA model performance are considered (i.e. locational agreement and landscape pattern structure), and it allows for manual intervention to set neighbourhood weighting parameters at any point. The remainder of this section presents the details of each stage.

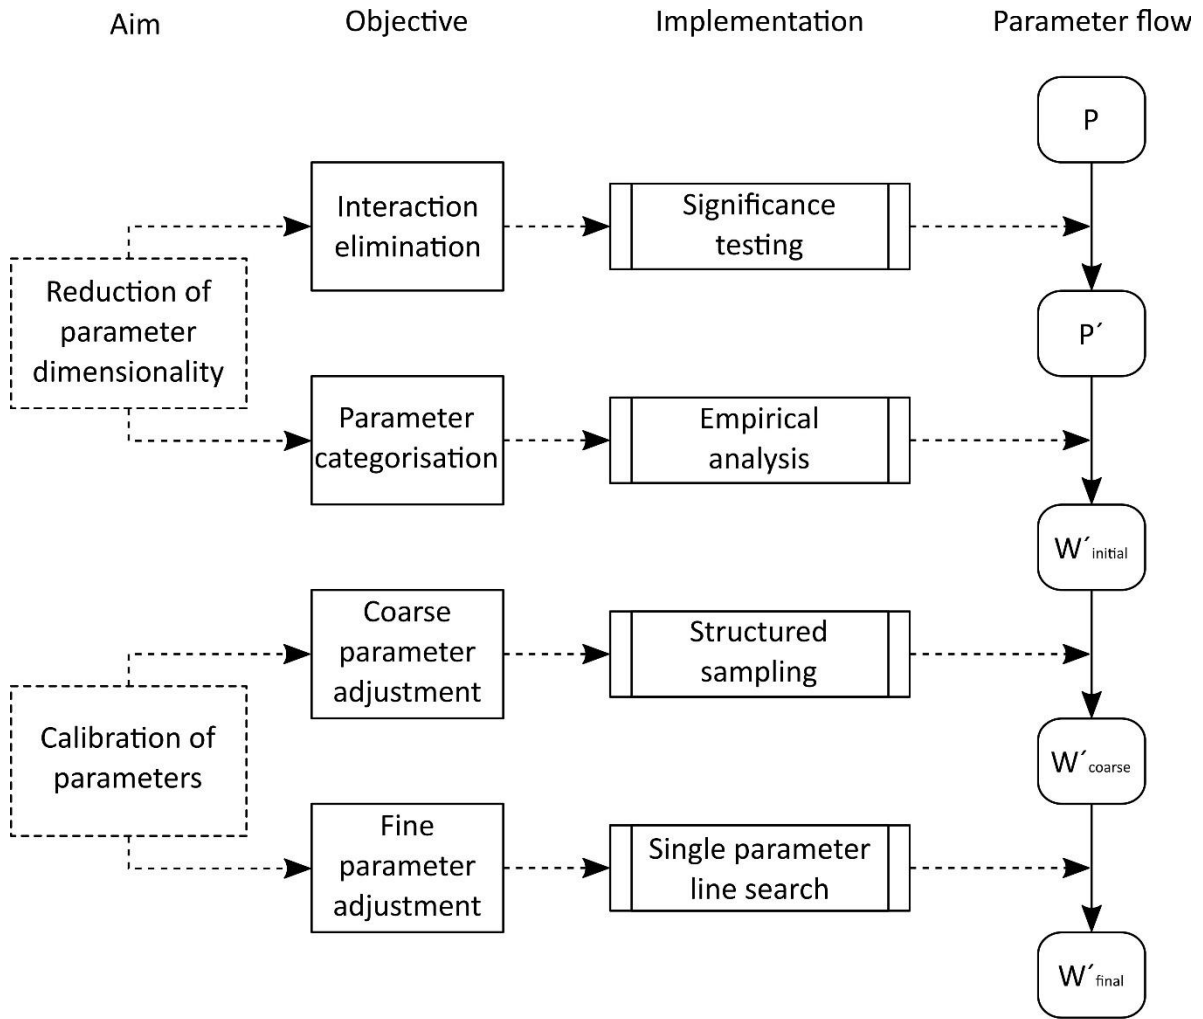


Figure 3.3. Conceptual outline of proposed methodology for semi-automatic process-specific neighbourhood rule calibration

### 3.3.1 Interaction elimination

#### 3.3.1.1 Overview

The interaction elimination stage reduces the parameter dimensionality by eliminating certain interactions to simplify the subsequent calibration problem, reducing the total set  $P$  to the smaller set  $P'$  of important interactions that are identified to be driving land-use changes in the region (Figure 3.3). The interaction elimination is based on an analysis of the available data, making the process objective and repeatable.

The primary objective of this stage is to determine the set of interactions  $P'$  composed of:

$$P' = P'_{IP} \cup P'_{CP} \cup P'_{ST} \cup P'_{CT} \quad (3.11)$$

where the subsets  $P'_{IP}$ ,  $P'_{CP}$ ,  $P'_{ST}$ , and  $P'_{CT}$  contain only the meaningful neighbourhood interactions. Typically, all inertia point and self-influence tail interactions are considered

meaningful, due to all land-uses exhibiting some tendency for persistence and some tendency for self-influence in the neighbourhood. This means that  $P'_{IP}$  is equal to  $P_{IP}$  and  $P'_{ST}$  is equal to  $P_{ST}$ , and that the reductions in  $P'$  are achieved in the reduction of conversion point and cross-influence tail interactions  $P'_{CP}$  and  $P'_{CT}$ .

The interaction elimination process is summarised in Figure 3.4, where an empirical analysis and significance test are used for identifying the meaningful interactions in  $P'_{CP}$  and  $P'_{CT}$ . This is based on an empirical evaluation of the transitions that occur in the calibration data  $\hat{\mathbf{X}}_0$  and  $\hat{\mathbf{X}}_d$ , as these data provide the main behaviour the land-use model is attempting to capture (Van Vliet et al., 2013b). The empirical analysis and significance test are conducted to determine if there is a statistically significant representation of a given land-use class for transitions to each active land-use class. This analysis is undertaken either at the location of the transitioned cells, informing the inclusion of a conversion point, or in the neighbourhood of the transitioned cells, informing the inclusion of a cross-influence tail.

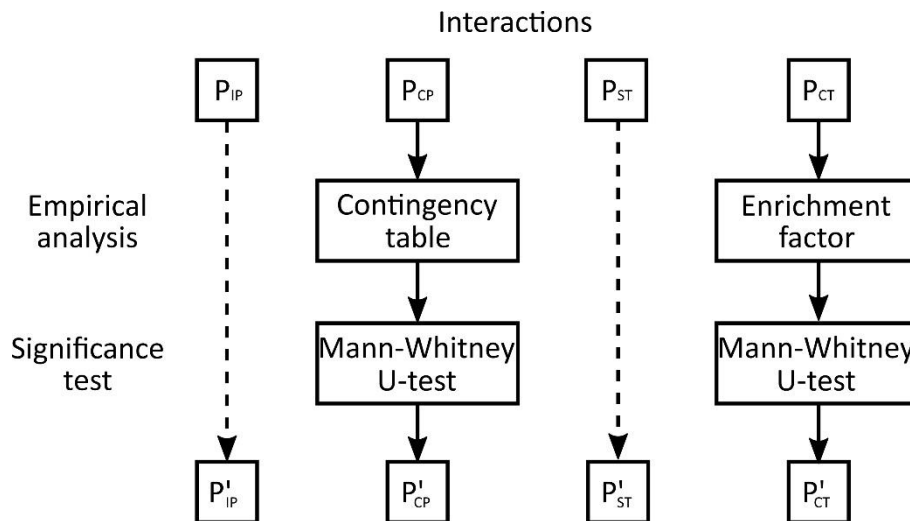


Figure 3.4. Proposed interaction elimination, based on empirical evaluation and significance testing

Hence, the meaningful interactions are such that:

$$\begin{aligned} P'_{CP} &= \{(i, j)_0 \in P_{CP} : T_0((i, j)_0, \hat{\mathbf{X}}) > t_0 \text{ and } Z_0((i, j)_0, \hat{\mathbf{X}}) > z_0\} \\ P'_{CT} &= \{(i, j)_1 \in P_{CT} : T_1((i, j)_1, \hat{\mathbf{X}}) > t_1 \text{ and } Z_1((i, j)_1, \hat{\mathbf{X}}) > z_1\} \end{aligned} \quad (3.12)$$

where  $T_d((i, j)_d, \hat{\mathbf{X}})$  is the empirically derived measure for the representation of land-use class  $j$  in the neighbourhood of cells at distance  $d$  that transitioned to class  $i$  (see Figure 3.4),  $t_d$  is the associated empirical threshold value,  $Z_d((i, j)_d, \hat{\mathbf{X}})$  is the measure of statistical significance for the representation of land-use class  $j$  in the neighbourhood of cells at distance  $d$  that transitioned

to class  $i$ , and  $z_d$  is the associated significance level. Interactions that do not return a neighbourhood representation that is above a user-defined threshold for  $t_d$  and  $z_d$  are eliminated from subsequent calibration.

The proposed interaction elimination is designed to extract the meaningful interactions based on the available data, which is likely limited to a pair of data maps. The method is inherently limited by the available data and can only capture what the available data indicate (i.e. intermediate transitions cannot be considered if there are no data to show these), which must be considered during application, as this has the potential to impact calibration performance (Blecic et al., 2015). The user must take into account the expected number of changes over the period considered when setting the thresholds. The user also has the ability to intervene following the interaction elimination, to either include rules that are not identified as statistically significant, or to eliminate rules that are identified as significant but are inconsistent with process understanding.

The empirical analysis measures and significance test used for the elimination are described in the next two sections. It should be noted that the elimination uses measures to capture the independent influence of individual classes on an active class, as this is consistent with the model structure (i.e. neighbourhood rules only capture individual influence of one land-use on another). So, for example, because the model uses a neighbourhood rule describing the influence of land-use class  $j$  on the allocation of class  $i$ , a measure relating the presence of land-use class  $j$  in the neighbourhood of class  $i$  is used for evaluation. Higher order, combined spatial influences are not considered, as these are not sufficiently captured by the transition potential model, with such effects only accounted for by the summation of the individual influences mentioned (see Equation 3.6).

### 3.3.1.2 Empirical analysis measures

The evaluation of conversion points is based on a form of the confusion matrix (Congalton, 1991), referred to as a contingency table (see Table 3.1 for an example), which is generated by logging the land-use class transitions for each cell for  $\hat{\mathbf{X}}_0$  to  $\hat{\mathbf{X}}_a$  (see Figure 3.4).

Table 3.1. Example contingency table, populated by logging the land-use class in each cell between two maps

		$\widehat{X}_a$			
		$A_1$	$A_2$	...	$A_n$
$\widehat{X}_0$	$A_1$	$\eta_{1,1}$	$\eta_{1,2}$	...	$\eta_{1,n}$
	$A_2$	$\eta_{2,1}$	$\eta_{2,2}$	...	$\eta_{2,n}$
	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
	$A_n$	$\eta_{n,1}$	$\eta_{n,2}$	...	$\eta_{n,n}$

In Table 3.1,  $\eta_{i,j}$  is the total number of cells that are class  $i$  in  $\widehat{X}_0$  and have transitioned to class  $j$  in  $\widehat{X}_a$ . The value  $\eta_{i,i}$  (the case where  $i$  is equal to  $j$ ) shows the total number of cells that did not change land-use class between  $\widehat{X}_0$  and  $\widehat{X}_a$ , which provides information about the level of inertia. Using the contingency table, the inertia rate, which quantifies the tendency of a land-use class to persist, can be derived as:

$$IR_i = \frac{\eta_{i,i}}{\sum_{m=1}^n \eta_{i,m}} \quad (3.13)$$

Using the contingency table, the conversion rate, which quantifies the tendency of transitions to a land-use class as a function of all transitions to that class, can also be derived as follows:

$$CR_{i,j} = \frac{\eta_{i,j}}{(\sum_{m=1}^n \eta_{m,j}) - \eta_{i,i}} \quad (3.14)$$

For the proposed approach, the threshold function for the inclusion of a conversion point ( $T_0$ ) is the conversion rate, derived from the contingency table.

The evaluation of cross-influence tails is performed using the enrichment factor of transitioning cells (Van Vliet et al., 2013b), which expresses the over- or under-representation of a particular land-use class in the neighbourhood of cells that transitioned to a certain class at a certain distance, relative to the representation of the neighbourhood class in the entire landscape (Figure 3.4):

$$EF_{i,j,d} = \log_{10} \left( \frac{\bar{R}_{i,j,d}}{\frac{n_j}{n}} \right) \quad (3.15)$$

where  $EF_{i,j,d}$  is the enrichment factor for the presence of land-use class  $j$  at distance  $d$  in the neighbourhood of cells that transitioned to class  $i$ ,  $\bar{R}_{i,j,d}$  is the average relative representation of land-use class  $j$  in the neighbourhood at distance  $d$  of cells that transitioned to land-use class  $i$ ,  $n_j$  is the total number of cells of land-use class  $j$  in the data map  $\widehat{X}_0$ , and  $n$  is the total number



of cells in the data map  $\hat{X}_0$ . To assist with interpretation, the enrichment factor is log-scaled by a factor of 10, so that values greater than 0 indicate over-representation, and values less than 0 indicate under-representation. The threshold function ( $T_I$ ) for the inclusion of a cross-influence tail is based on the log-scaled enrichment factor at a distance  $d$  equal to one.

It should be noted that the information from the contingency table, specifically the conversion rate, is used for the analysis of conversion points as opposed to the enrichment factor values at distance zero because the contingency table is more effective at capturing the different conversions in the data. Although the two measures use the same information (the enrichment factor at distance zero can be derived from the information in the contingency table, see Supplementary material 3A), the enrichment factor can be less effective at capturing the conversions occurring. For example, many conversions may be occurring from a certain class to another, but the enrichment factor may not suggest over-representation because there is a large representation of the class in the landscape (large  $n_i$ ), erroneously indicating that a conversion point is not required. A common example of this are conversions from large, relatively passive land-use classes, such as natural vegetation, because they occupy a large area of the landscape, and facilitate many conversions to different classes, which are meaningful to include.

### 3.3.1.3 Significance test

The Mann-Whitney U-test (MWU-test) is applied to determine whether the representation of land-use class  $j$  in transitions to land-use class  $i$  is statistically significant (Figure 3.4). The MWU-test requires the compilation of ranked data sets (shown in Figure 3.5 for the neighbourhood at distance one), for (i) the percentage of cells of land-use class  $j$  in the neighbourhood at distance  $d$  to cells transitioning to class  $i$ , compared to (ii) the percentage of cells of land-use class  $j$  at distance  $d$  for all cells. The MWU-test is used to assess whether there is significant variation for data sets (i) and (ii). This is formally stated as testing whether a sample drawn from one distribution is equally likely to be greater or less than a sample drawn from the other (this is the null hypothesis,  $H_0$ , where the alternative hypothesis,  $H_A$ , is that the first aforementioned sampled is more likely to be either greater or less than the second). For the proposed approach,  $H_0$  is that the ranks of the respective relative composition tallies for data sets (i) and (ii) are statistically equivalent. This is shown in Figure 3.5 for case *a*, as the distributions appear fairly similar. Alternatively,  $H_A$  is that the ranks of the respective relative composition tallies for the two data sets are systematically higher or lower. This is shown in Figure 3.5 for case *b*, as the distributions for the transition data set possess a significantly larger

percentage of higher cell counts (i.e. ~15% of transition cells have 8 cells of class  $j$  at distance  $d$  equal to 1, while only ~2% of all cells have this number at the same distance).




Figure 3.5. An example tally of the frequency of cells of class  $j$  at distance  $d$  equal to 1 from a given land-use class transition (bottom row), as compared to the frequency of cells of class  $j$  at distance  $d$  equal to 1 from all cells (top row). For case a, the distributions appear similar, whereas for case b, the distributions are skewed to different extremes, suggesting differences that are statistically significant.

The MWU-test requires the U-test statistic to be calculated, which requires the two distribution tallies to be combined and ranked from low to high (Corder and Foreman, 2014). With the ranked order determined, the U-test statistic is calculated for the ranks of the distribution of a specific land-use class  $j$  in the neighbourhood at a distance  $d$  of cells that transitioned to land-use class  $i$ :

$$U_{i,j,d} = n_i \cdot n + \frac{n_i(n_i + 1)}{2} - RS_{i,j,d} \quad (3.16)$$

where  $U_{i,j,d}$  is the U-test statistic for the presence of land-use class  $j$  at a distance  $d$  from cells that transitioned to land-use class  $i$ ,  $n_i$  is the number of newly allocated cells of class  $i$ ,  $n$  is the total number of cells in the landscape, and  $RS_{i,j,d}$  is the sum of the ranks of the distribution of the relative composition for the presence of land-use class  $j$  at a distance  $d$  from cells that

transitioned to land-use class  $i$ . The U-test statistic is also calculated for the ranks of the relative composition of a specific land-use class  $j$  at a distance  $d$  for all cells:

$$U_{c,j,d} = n_i \cdot n + \frac{n(n+1)}{2} - RS_{c,j,d} \quad (3.17)$$

where  $U_{c,j,d}$  is the U-test statistic for the presence of land-use class  $j$  at a distance  $d$  of all cells  $C$ , and  $RS_{c,j,d}$  is the sum of the ranks of the distribution of the relative composition for the presence of land-use class  $j$  at a distance  $d$  of all cells  $C$ . The final U-test statistic used is the minimum of the two that are calculated.

For this proposed application, the samples are sufficiently large that a normal approximation can be generated for the determination of significance, which corresponds to  $Z_d$  in Equation 3.12:

$$Z(i,j)_d = \frac{\min\{U_{i,j,d}, U_{c,j,d}\} - \bar{x}_U}{s_U} \quad (3.18)$$

where  $Z_d$  is the z-score approximation for the presence of land-use class  $j$  at a distance  $d$  of cells that transitioned to land-use class  $i$ ,  $\bar{x}_U$  is the mean of the data, and  $s_U$  is the standard deviation of the data.

### 3.3.2 Parameter categorisation and initialisation

The parameter categorisation and initialisation stage further reduces the dimensionality of the calibration problem to facilitate initialising the LUCA model with a set of neighbourhood rule weighting parameter values for subsequent coarse calibration (Figure 3.3) by using a method to efficiently generate a set of neighbourhood rules that are consistent with process understanding. The categorisation process achieves this across two steps, discussed in detail below.

The first step is the introduction of a set of meta-parameters,  $\theta_{CP}^0$ ,  $\theta_{ST}^0$ , and  $\theta_{CT}^0$ , which describe the *inter-type* importance of each interaction type  $P'_{CP}$ ,  $P'_{ST}$ , and  $P'_{CT}$  with respect to the inertia point type interaction type  $P'_{IP}$ . The meta-parameters quantify how important each interaction type is relative to the inertia point interaction type, which is assumed to be the dominant process (e.g. the meta-parameter quantifies how important the  $P'_{CP}$  interactions are compared to the  $P'_{IP}$  interactions). For example, high values of  $\theta_{CP}$  mean conversion points exhibit similar influence to inertia points, which will likely cause more conversion. The meta-parameters range from zero to one, and are set by the user (note that a value of 1 places the interactions at the same

level of importance as the inertia point interactions). Recommended ranges are between 0 and 0.1 for all meta-parameters.

The second step is the categorisation of the parameters included for calibration within each interaction type. This achieves a discretisation of the corresponding weighted influence values ( $w_{i,j,d}$ ), where rather than taking values from a continuum, they are restricted to a set of finite values. This categorisation establishes a hierarchy of the *intra-type* importance of each interaction for each type of interaction (e.g. how important interaction  $(i,j)_0$  is compared to  $(i,k)_0$  for both  $(i,j)_0, (i,k)_0 \in P'_{CP}$ ). Categorisation is based on an empirical analysis of the available data.

A result of the proposed categorisation method is that the weighting parameters for the interactions  $P'_{IP}, P'_{CP}, P'_{ST}$ , and  $P'_{CT}$  are given by the following representations:

$$\begin{cases} w_{i,i,0} = \tilde{w}_{i,i,0} \\ w_{i,j,0} = \theta_{CP}^0 \cdot \tilde{w}_{i,j,0} \\ w_{i,i,d} = \theta_{ST}^0 \cdot \tilde{w}_{i,i,1} \cdot u(d) \\ w_{i,j,d} = \theta_{CT}^0 \cdot \tilde{w}_{i,j,1} \cdot u(d) \end{cases} \quad (3.19)$$

where  $\theta_{CP}^0, \theta_{ST}^0$ , and  $\theta_{CT}^0$  are the meta-parameters that are heuristically assigned and take values between 0 and 1, and  $\tilde{w}_{i,j,d}$  are the normalised weighting parameters that take values from the discretised parameter space  $\{\tilde{w}_{i,j,d,[1]}, \dots, \tilde{w}_{i,j,d,[K]}\}$ , where  $K$  is the number of discrete levels. So, for example, if an interaction type has three levels, low, medium and high, the weights take values from the set range  $\{w_{i,j,d,[1]}, w_{i,j,d,[2]}, w_{i,j,d,[3]}\}$ .

The parameter categorisation is based on an empirical evaluation of the available data. A point or tail for each interaction type is assigned as follows:

$$\tilde{w}_{i,j,d} = \begin{cases} \tilde{w}_{i,j,d,[1]} & \text{if } T_k((i,j)_d, \hat{\mathbf{X}}) \leq I_k^1 \\ \tilde{w}_{i,j,d,[2]} & \text{if } I_k^1 < T_k((i,j)_d, \hat{\mathbf{X}}) \leq I_k^2 \\ \vdots & \vdots \\ \tilde{w}_{i,j,d,[K]} & \text{if } I_k^{K-1} \leq T_k((i,j)_d, \hat{\mathbf{X}}) \end{cases} \quad (3.20)$$

where  $T_k$  is the empirically determined threshold value for interaction type  $k$  (i.e.  $IP, CP, ST, CT$ ),  $d$  is either 0 or 1, defining either a point or a tail, and  $I_k^n$  is the upper threshold value for importance category  $n$  (different for each interaction type, where higher  $n$  means higher importance).

The categorisation of the different interaction types uses the measures previously derived in Section 3.3.1.2, shown in Table 3.2. The  $P'_{IP}$  and  $P'_{CP}$  interactions are categorised using

measures derived from the contingency table, the inertia rate and conversion rate, respectively. The  $P'_{ST}$  and  $P'_{CT}$  interactions are categorised using the enrichment factor values at distance  $d$  equal to one. Thresholds ( $I_k$ ) for the different groups are set by the user.

*Table 3.2. Summary of empirical measures used to categorise different interaction types for the proposed categorisation*

<b>Interaction type</b>	<b>Empirical measure</b>	<b>Functional form</b>	<b>Equation</b>
Inertia point	Inertia rate	$IR_i$	3.13
Conversion point	Conversion rate	$CR_{i,j}$	3.14
Self-influence tail	Enrichment factor	$EF_{i,i,l}$	3.15
Cross-influence tail	Enrichment factor	$EF_{i,j,l}$	3.15

Hence, given that each interaction type uses the same number of levels, the calibration complexity has been reduced so that the LUCA model neighbourhood rules are characterised by three meta-parameters, and weighting parameter values for each discrete level across all interaction types. The model can be initialised by specifying these values, and the required thresholds, to generate an initial parameter estimate  $W'_{initial}$ .

### **3.3.3 Coarse parameter adjustment**

The coarse parameter adjustment stage calibrates the neighbourhood weighting parameters at a coarse level to improve LUCA model performance by tuning the categorised neighbourhood weighting parameters introduced in Section 3.3.2 (Figure 3.3). The coarse parameter adjustment facilitates the initialisation of the fine parameter adjustment process at a starting point within the parameter space that results in objectively good performance and is consistent with process understanding. Model performance is measured by comparing the simulated output with the data using at least two metrics, as discussed in the Section 3.1, one to quantify locational agreement and another to quantify landscape pattern structure. Model performance is improved by generating a set of parameters  $W_{CP}^{coarse}$ ,  $W_{ST}^{coarse}$ , and  $W_{CT}^{coarse}$  through the coarse adjustment of the meta-parameters  $\theta_{CP}$ ,  $\theta_{ST}$ , and  $\theta_{CT}$ . Within this stage, the parametric representation (Equation 3.19) is retained, but the meta-parameters are varied about the initial settings.

By using the neighbourhood weighting parameterisation from Equation 3.19, the weights can be expressed as functions of the meta-parameters such that:

$$\begin{cases} w_{CP} = w_{CP}(\theta_{CP}) \\ w_{ST} = w_{ST}(\theta_{IT}) \\ w_{CT} = w_{CT}(\theta_{CT}) \end{cases} \quad (3.21)$$

Using this formulation, the initial parameter set is given by:

$$\begin{cases} w_{CP}^0 = w_{CP}(\theta_{CP}^0) \\ w_{ST}^0 = w_{ST}(\theta_{ST}^0) \\ w_{CT}^0 = w_{CT}(\theta_{CT}^0) \end{cases} \quad (3.22)$$

As outlined above, the coarse level parameter adjustment is performed by adjusting the meta-parameters to optimise the objectives of calibration performance. Formally, the goal is to identify a set of meta-parameters  $\theta^* = \{\theta_{CP}^*, \theta_{ST}^*, \theta_{CT}^*\}$ , that result in non-dominated calibration performance (i.e. improved performance in one objective cannot be achieved without declined performance in the other objective) to the bi-objective problem:

$$\begin{aligned} \max f_1(W'(\theta), \hat{X}_a) \\ \min f_2(W'(\theta), \hat{X}_a) \end{aligned} \quad (3.23)$$

where  $f_1$  corresponds to the locational agreement objective between the calibration data map  $\hat{X}_a$  and the LUCA model simulation map  $X_a$ , which is a function of  $W'(\theta)$ , and  $f_2$  corresponds to the landscape pattern structure objective, measured as the error between the calibration data map  $\hat{X}_a$  pattern, and the LUCA model simulation map  $X_a$  pattern, which is a function of  $W'(\theta)$ . As opposed to determining every set of meta-parameters that optimises the objective, an approximate (and more computationally efficient approach) is implemented as shown in Figure 3.6. In this approach, one meta-parameter is varied at a time, within each stage, allowing for an explicit consideration of the non-dominated solutions subject to the variation of the single meta-parameter under consideration. Within this process, user input is required to select the best meta-parameter value from the set of solutions. This is an important step in the process as it allows the user to make a judgement as to which meta-parameter value is best based on additional subjective criteria (e.g. user experience and/or consideration of landscape features not characterised by the objective metrics). The selected meta-parameter value, as shown in Figure 3.6, is then used within the subsequent stages, where the other meta-parameters are varied one at a time. The final set of meta-parameters  $\theta^*$  is obtained at the conclusion of the coarse adjustment stage.

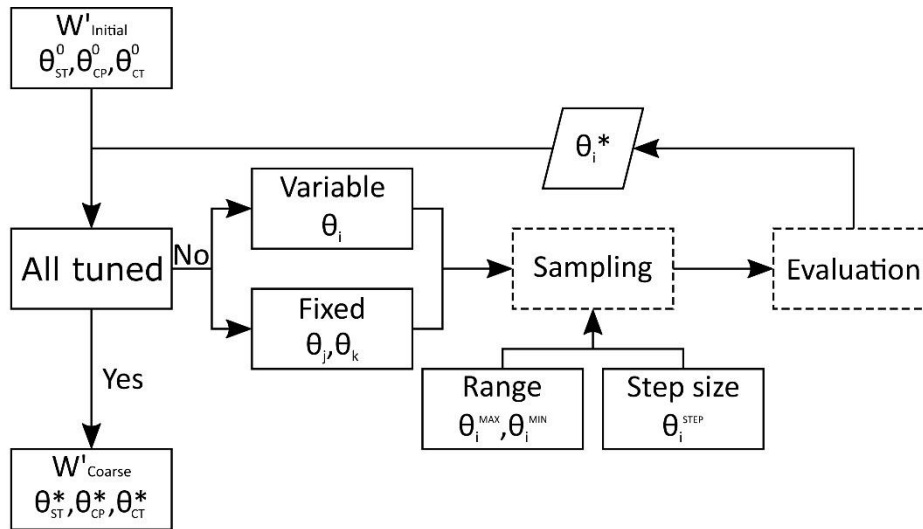


Figure 3.6. Iterative procedure for performing coarse parameter adjustment where meta-parameters are tuned.

### 3.3.4 Fine parameter adjustment

The fine parameter adjustment further calibrates the neighbourhood weighting parameters, refining those obtained after the coarse adjustment stage to generate a final set of calibrated neighbourhood weighting parameters for subsequent validation and future scenario analysis (Figure 3.3). The fine parameter adjustment stage is initialised using the output of the coarse adjustment stage, as this corresponds to a good starting point in the parameter space that has objectively good performance and agreement with process understanding. As with the coarse level adjustment, locational agreement and landscape pattern structure metrics are used to objectively assess calibration performance.

The fine parameter adjustment stage considers the neighbourhood weighting parameters on an individual level, rather than the meta-parameters, as was the case in the coarse parameter adjustment stage, iteratively refining each to optimise the performance metrics used, within the available number of iterations (computational budget). This is achieved using a line-search algorithm, as shown in Figure 3.7, because this is an automatic search method that efficiently converges on a parameter value that locally optimises the objectives. The specific neighbourhood weighting parameter  $w_{i,j,d}$ , corresponding to one of the four interactions groups, is assumed to lie within a range for which a certain value optimises the performance metrics used. The line-search is conducted using the Golden-section search algorithm.

The over-arching structure of the sequential line-search algorithm is as follows. The working parameter set  $W^i$  is initialised to the values obtained at the conclusion of the coarse calibration,

as shown. Then, the algorithm loops through the interaction types (in order of importance that is user defined and case study dependent, generally inertia points, self-influence tails, conversion points and cross-influence tails). For each type, the algorithm sequentially loops through all weighting parameters, searching between a set of user-defined minimum and maximum neighbourhood weighting parameter values  $w_{i,j,d}^{Min}$  and  $w_{i,j,d}^{Max}$  to determine the refined value  $w_{i,j,d}^{fine}$  that maximises the specified objective  $f_{fine}$  (discussed below), subject to all other parameters being held constant. The working parameter set is updated when the refined variable value is obtained, and the procedure continues through the loop until the computational budget is exhausted, returning the final set,  $W'_{final}$ , composed of the refined neighbourhood weighting parameters.

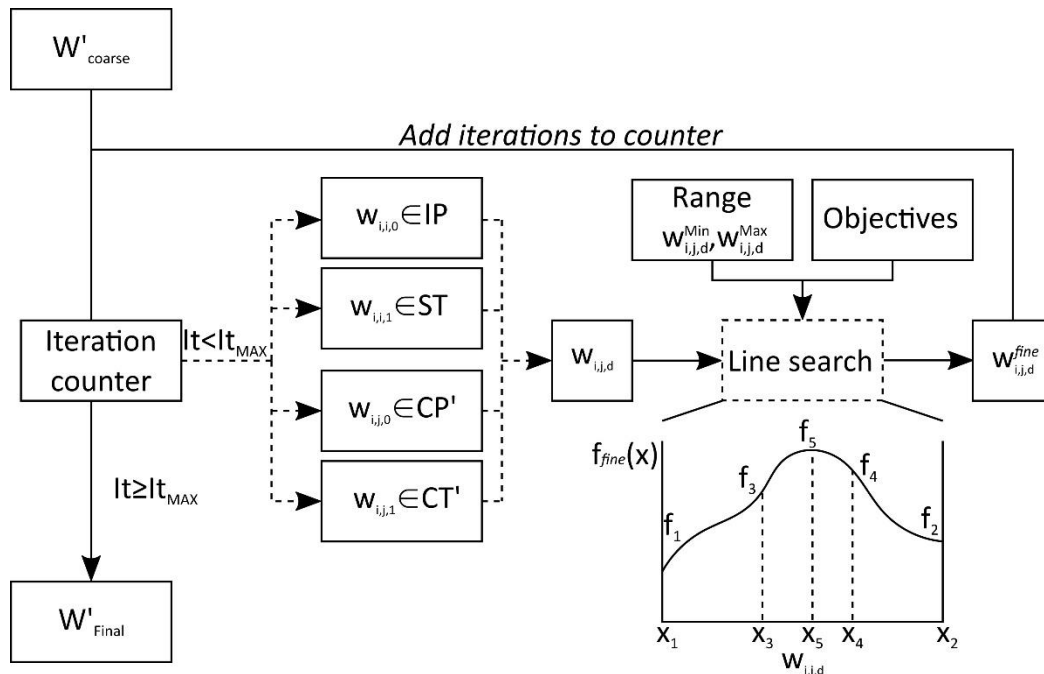


Figure 3.7. Iterative procedure for performing fine parameter adjustment via a line-search algorithm applied to individual neighbourhood weighting parameters

The line-search uses a single objective,  $f_{fine}$ , to characterise the calibration performance of the proposed parameter set. As discussed previously, multiple objectives are required to characterise the performance of LUCA models, so  $f_{fine}$  is taken as the weighted sum of the objectives:

$$f_{fine}(W', \hat{X}) = \sum_{i=1}^m a_i f_i(W', \hat{X}) \quad (3.24)$$



where  $f_i$  is the  $i^{\text{th}}$  objective,  $a_i$  is the user-defined weight given to the  $i^{\text{th}}$  objective (all weight values sum to one), and  $m$  is the number of objectives. As  $f_{\text{fine}}$  is the sum over both locational agreement and landscape pattern structure error, each metric must be transformed into a maximisation objective (e.g. reverse the sign of the metric of minimisation objectives). Additionally, to mitigate potential issues with variable ranges of the metrics, the objectives are scaled to values between 0 and 1, using metric ranges as the normalising upper and lower bounds that are defined by the user. It is important that the metrics used are appropriately balanced to achieve a trade-off between the performance objectives to ensure robust calibration and to prevent over-calibration.

The proposed approach assumes a model structure sufficiently general to capture the major land-use change processes in a region. This assumption allows for a focus on efficient calibration, where all objectives can be combined using a single weighted sum. However, this does limit the ability to fully explore the trade-off between the objectives, which can obscure more fundamental modelling issues. A notable example would be if a balance between locational agreement and landscape pattern structure cannot be easily achieved (White et al., 2015), where the calibrated model performs well in either metric, but not both. This is important to consider, as it might suggest more fundamental structural model issues, such that the LUCA model does not fully capture major land-use change processes. Identifying and addressing such structural problems are not a primary focus of this work, but this is an important issue to consider that warrants future research.

### **3.4 Application of proposed approach**

The utility of the proposed semi-automatic calibration method for neighbourhood rules is evaluated using the following computational testing regime. The approach is applied to calibrate the neighbourhood rules of four case studies in Europe, comparing the output obtained based on the preference for different objectives, and evaluating the performance against neutral models of landscape change that are used to generate benchmark metric values for calibration performance (Hagen-Zanker and Lajoie, 2008).

#### **3.4.1 Land-use model**

The LUCA model Metronamica is used to model land-use changes. Metronamica is a generic, constrained CA modelling framework (Van Delden and Hurkens, 2011) that has numerous global applications (Wickramasuriya et al., 2009, Rutledge et al., 2008, Van Delden et al., 2011). Metronamica is derived from the transition potential model developed by White and

Engelen (1993c), and hence includes neighbourhood rules in the determination of land-use changes. Metronamica follows the convention for the calculation of transition potential (Equation 3.4), where the additional processes considered via parametric maps are given by:

$$\theta_{c,i} = A_{c,i} \cdot S_{c,i} \cdot Z_{c,i} \quad (3.25)$$

where  $A_{c,i}$  is the influence of accessibility,  $S_{c,i}$  is the influence of suitability, and  $Z_{c,i}$  is the influence of zoning. The specific parameterisation of each component is given in the Metronamica documentation (RIKS, 2015).

### **3.4.2 Case studies**

The four case studies used for testing include Berlin, Germany; Budapest, Hungary; Lisbon, Portugal; and Madrid, Spain, and their surrounding regions, shown in Figure 3.8. These case studies are selected to enable the approach to be tested under a range of different conditions, given the variation in physical characteristics (e.g. coastal or inland, high or low elevation) and rates of urban growth (e.g. low for Berlin, moderate for Budapest, and high for Lisbon and Madrid, see Supplementary material 3C for contingency tables) of these locations.

The CORINE land-use data set (Haines-Young et al., 2006) is used to generate the land-use maps for the case studies. Each case study uses 15 land-use classes, eight of which are actively modelled, which are reclassified from the 48 CORINE level 3 land-use classes, with a 250-metre resolution covering a region of 10,000 km<sup>2</sup> (400 by 400 cells), centrally located on the city centre. The calibration period for each case study is 1990-2000 and the validation period is 2000-2006. Each case study includes major road data for accessibility.

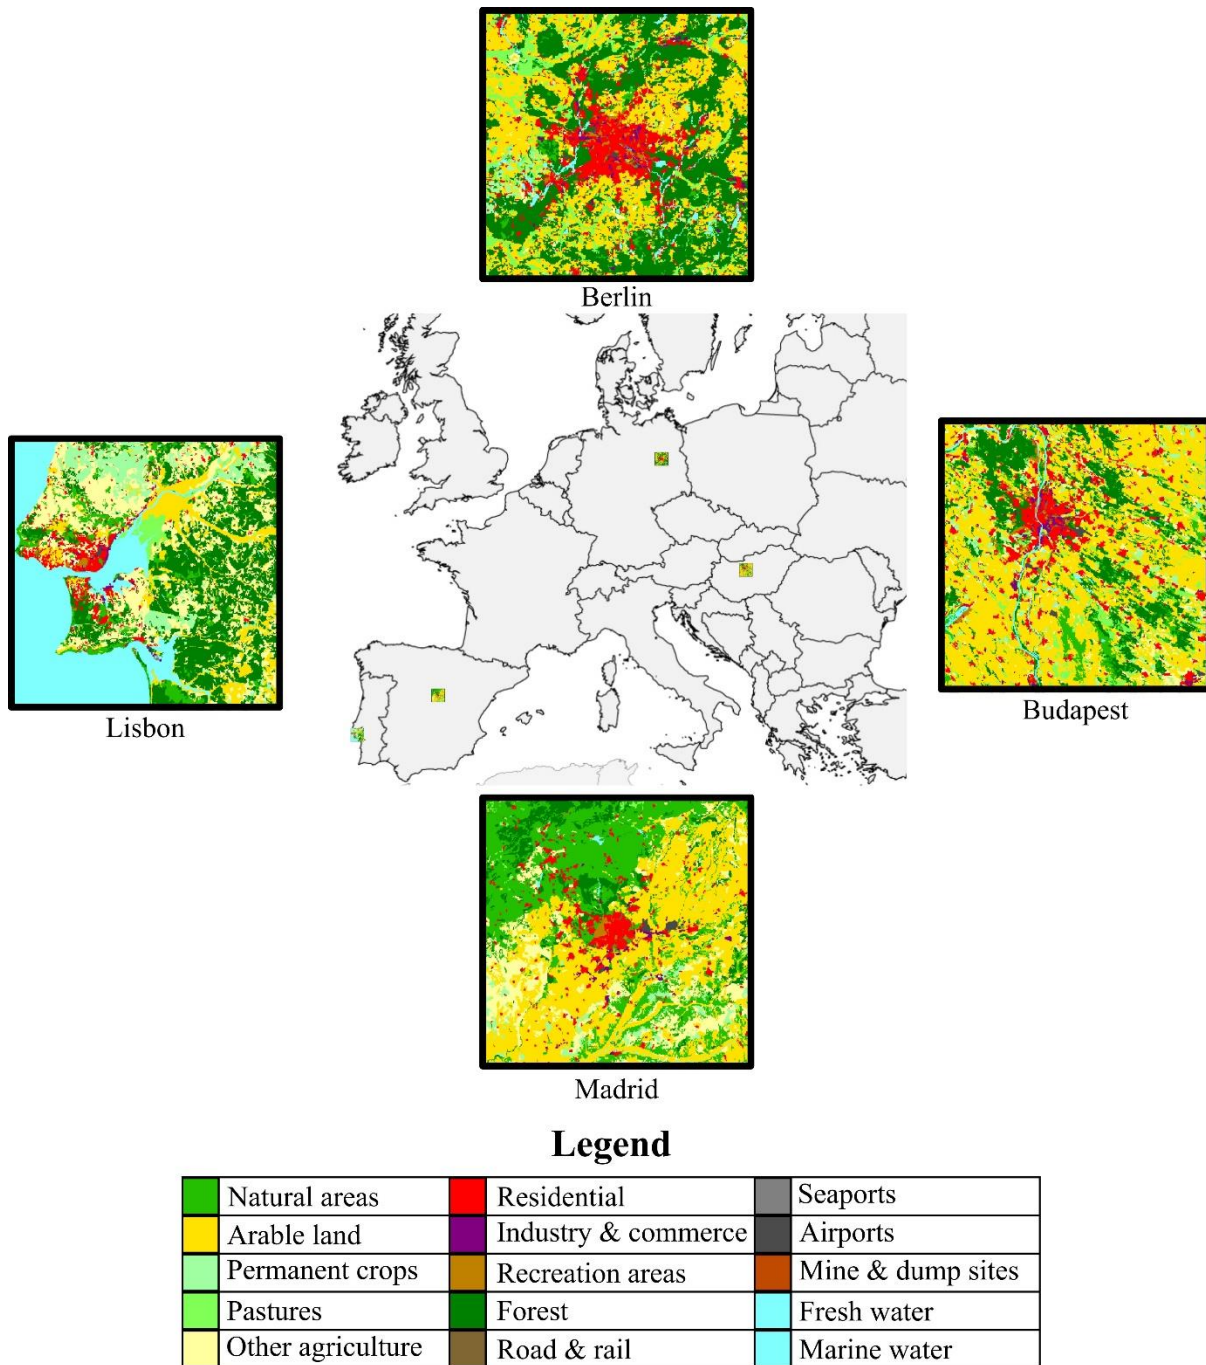


Figure 3.8. 1990 land-use maps of the four European case studies.

### 3.4.3 Implementation of proposed methodology

For each of the case studies, the Empirical Neighbourhood rule Calibrator (ENC) software developed as part of this work is applied. As each case study uses the same number of total and active land-use classes, each is implemented with the same number of possible neighbourhood interactions and possible neighbourhood weighting parameters.

### 3.4.3.1 Interaction elimination

To conduct the interaction elimination stage, the thresholds presented in Table 3.3 are used for the empirical neighbourhood analysis. Only conversion points and cross-influence tails are eliminated for this proposed implementation, because all eight active classes for each case study are expected to exhibit inertia and self-influence. A conversion rate of 2.5% is used as a minimum threshold as this prevented infrequent and erroneous conversions present in the data, such as conversions from *fresh water* to *residential* land-use, from being included in the subsequent calibration. To demonstrate the most straight-forward application possible, the interaction elimination is applied to only identify attractive cross-influence tails, and does not consider repulsive cross-influences. Log-scaled enrichment factor values at distance one that indicated over-representation (i.e. greater than 0), and hence hint at a possible attractive influence between different land-uses, are used for the empirical analysis of cross-influence tails. A minimum z-limit of 1.96, the 95% confidence limit, is used as the statistical significance level for both conversions points and cross-influence tails.

Table 3.3. Thresholds used for inclusion in calibration

<b>Interaction type</b>	$T_d$	$Z_d$
Conversion point	$CR_{i,j} > 2.5 \%$	$Z_{i,j,0} > 1.96$
Cross-influence tail	$\text{Log}_{10}(EF_{i,j,1}) > 0$	$Z_{i,j,1} > 1.96$

The results of the parameter elimination are summarised in Figure 3.9, which shows the parameter reduction per case study for each interaction type. As shown, the implemented parameter elimination for the conversion points and cross-influence tails resulted in a substantial reduction in the number of interactions considered for each case study. The interactions that are included were generally consistent with process knowledge. For example, it is expected that attractive influences exist between the urban classes *residential*, *industry & commerce*, and *recreation areas*, given the limited number of urban classes. However, there are also examples, such as conversion points from agricultural land-uses to urban land-uses, which represent a case where the land-use is desirable for urban land-uses because it acts as a supply of vacant land, rather than being specifically attractive based on the agricultural land-use class. Such examples highlight the potential benefits of making the calibration procedure semi-automatic, as discursive knowledge can identify such parameters and remove them from subsequent calibration, or ensure these influences are minimised by the calibration procedure.




Figure 3.9. Number of interactions considered for conversion points and cross-influence tails before and after the interaction elimination stage for each case study

### 3.4.3.2 Parameter categorisation

To conduct the parameter categorisation stage, three possible categories, including low, medium, and high, are used for each case study. This allowed for sufficient differentiation of weighting parameter values for each interaction type. The threshold values for the categories used for each type of interaction are the same for each case study, and are presented in Table 3.4. Interactions not exceeding the medium threshold are graded as low. The values used are based on knowledge of the case studies, and some trial-and-error analysis to ensure sufficient category variation across the case studies.

Table 3.4. Thresholds used for categorisation of parameters as high, medium or low for each interaction type

Interaction type	$I_k^{High}$	$I_k^{Med}$	$I_k^{Low}$
Inertia point	$IR_i \geq 95\%$	$IR_i \geq 90\%$	$IR_i < 90\%$
Conversion point	$CR_{i,j} \geq 50\%$	$CR_{i,j} \geq 10\%$	$CR_{i,j} < 10\%$
Self-influence tail	$Log_{10}(EF_{i,i,1}) \geq 1.0$	$1.0 > Log_{10}(EF_{i,i,1}) \geq 0.5$	$Log_{10}(EF_{i,i,1}) < 0.5$
Cross-influence tail	$Log_{10}(EF_{i,j,1}) \geq 1.0$	$1.0 > Log_{10}(EF_{i,j,1}) \geq 0.5$	$Log_{10}(EF_{i,j,1}) < 0.5$

At this stage, a neighbourhood shape function that describes the distance-decay of the neighbourhood influence is used to aggregate self-influence and cross-influence tail neighbourhood weighting parameters. To make the calibration as efficient as possible, an aggregation strategy is used that aggregates tails to a set of key points (White et al., 1997), an influence value at distance one, an influence value at distance two that is ten percent of the influence at distance one, and a point at distance five where influence is set to zero. Using this strategy, each self-influence and cross-influence tail is represented by a single weighting parameter at distance one, reducing the complexity of the calibration problem.

To generate the initial neighbourhood weighting parameters, the inertia point parameters are set such that high inertia is assumed to have double the influence of medium inertia, and medium inertia is assumed to have double the influence of low inertia, using values of 1,000, 500 and 250 for high, medium and low inertia, respectively. The initial meta-parameters used are presented in Table 3.5. The values used are based on the case studies, which are mostly characterised by persistence, resulting in the highest values for the self-influence tail meta-parameter and lower values for the conversion point and cross-influence tail meta-parameters.

#### 3.4.3.3 Coarse parameter adjustment

The coarse parameter adjustment stage is conducted using the ranges and step-sizes for each meta-parameter given in Table 3.5. The order of the meta-parameter sampling is  $\theta_{ST}$ ,  $\theta_{CP}$ , and  $\theta_{CT}$ , ordered by the impact (i.e. value range) on the output metrics obtained. The ranges are determined based on discursive knowledge and some trial-and-error analysis, with values outside the ranges given in Table 3.5 resulting in poor performance in all objectives.

Table 3.5. Ranges and step sizes tested for meta-parameter sampling

Meta-parameter	Initial ( $\theta_k^0$ )	Minimum	Maximum	Step-size	Intervals
$\theta_{ST}$	0.050	0.00	0.10	0.005	20
$\theta_{CP}$	0.025	0.00	0.05	0.0025	20
$\theta_{CT}$	0.005	0.00	0.02	0.001	20

To evaluate locational agreement, two variations of Cohen's Kappa are used, Fuzzy Kappa (FK) developed by Hagen-Zanker (2009), and Fuzzy Kappa Simulation (FKS) developed by Van Vliet et al. (2013b). As both metrics are derivatives of Cohen's Kappa, they measure the observed agreement between two categorical data sets, corrected for the agreement expected from random allocation of the given class sizes, with fuzziness allowing for the consideration

of partial agreement of location and class. The major distinction between the two metrics is that FK considers the entire land-use map for the calculation of agreement, whereas FKS only considers the transitioned cells. Both metrics are used because the variation in the measurement of agreement can impact the results obtained (Newland et al., 2018a).

To measure landscape pattern structure, the error of the simulated and observed clumpiness for the actively modelled land-use classes is used. Clumpiness is a measure of the proportional deviation of the proportion of like adjacencies from that expected under a spatially random distribution for a specific land-use class (McGarigal, 2014). As clumpiness is measured at the class level, it requires aggregation to a single value. Previous calibration methods that used the average class level clumpiness (Newland et al., 2018a) found this is not an ideal aggregation strategy, because it can over-emphasize relatively minor land-use classes that occupy a relatively small amount of the total landscape. Hence, the absolute Area-Weighted Clumpiness Error (AWCE) of the actively modelled land-use classes is used, as this provides a better landscape-centric perspective by emphasizing the classes that occupy the greatest area within the region.

Given the potential uncertainty surrounding the impact of the different metrics on the performance of the proposed approach, three combinations of metrics are used, as shown in Table 3.6. Each combination includes a metric for locational agreement and landscape pattern structure. FK and FKS are included together in set three to provide further insight into how these objectives affect the performance of the proposed approach.

*Table 3.6. Combinations of metrics used as objectives for calibration*

<b>Set</b>	<b>Metrics</b>
1	FK, AWCE
2	FKS, AWCE
3	FK, FKS, AWCE

During the coarse adjustment stage, two types of behaviour are observed in the objective space, which dictates the selection of the meta-parameter value (Figure 3.10). The first type of behaviour is *convergent*, shown in plot *a* for the FKS and AWCE values for the range of  $\theta_{ST}$  values sampled for the Madrid case study. As shown, there is a meta-parameter value that produces both the best FKS and AWCE value, which makes selecting the meta-parameter value straightforward. The second type of behaviour observed is a *trade-off*, shown in plot *b* for the FKS and AWCE values for the range of  $\theta_{ST}$  values sampled for the Berlin case study. As shown,

improved performance in FKS results in reduced performance in AWCE (as the error increases). In this case, selection of the most appropriate meta-parameter requires further interpretation from the user, as the selection of a meta-parameter value requires a preference for a certain trade-off between objectives, which impacts the fine parameter adjustment starting position and hence the resultant final output.




Figure 3.10. Example of the different behaviour observed in the objective space for FKS and AWCE when tuning a meta-parameter: convergent (a), where the selected point optimises both metrics, and trade-off (b), where improved performance in one objective results in reduced performance in the other. A trade-off requires user interpretation to preference a certain trade-off between objectives.

#### 3.4.3.4 Fine parameter adjustment

The fine parameter adjustment stage of calibration uses the combinations of metrics listed in Table 3.6. The resultant output obtained from the fine parameter adjustment stage is heavily influenced by the weight and ranges used for each metric, allowing the user to preference a certain objective. An example is shown in Figure 3.11 for the trajectory of the metrics whilst conducting the fine parameter adjustment for the Lisbon case study using the objectives FK and AWCE. As shown, given the same starting point, a preference for either locational



agreement, landscape pattern structure, or a balance between the two, has a large impact on both the trajectory and the final metric values obtained.




Figure 3.11. Example of the different behaviour observed in the objective space for FK and AWCE during fine adjustment of the Lisbon case study, based on a preference for a certain objective. A preference tends to strongly bias the results to a certain metric, and a balanced preference can lead to improvement in both metrics.

Given the possible variation that can be achieved, three tests are conducted for each set of metrics for each case study in order to test the impact of the selection of this preference on the performance of the proposed algorithm: one test that purely focuses on improving Locational Agreement (LA), one test that purely focuses on improving Landscape Pattern Structure (LPS), and one test that balances improvements in both objectives. The weightings used for each test are summarised in Table 3.7. As shown, total preference for a specific objective is achieved by weighting the other objective(s) as zero, essentially making the problem single objective. This is to highlight potential issues with using a single objective, as this will likely result in over-calibration. It should be noted that FK and FKS are evenly weighted for the case where both are used and there is preference for locational agreement, or a balanced preference.

Table 3.7. Weighting used for each metric depending on objective preference

Metrics	Preference	FK weight	FKS weight	AWCE weight
FK, AWCE	LA	1.00	N/A	0.00
	Balanced	0.50	N/A	0.50
	LPS	0.00	N/A	1.00
FKS, AWCE	LA	N/A	1.00	0.00
	Balanced	N/A	0.50	0.50
	LPS	N/A	0.00	1.00
FK, FKS, AWCE	LA	0.50	0.50	0.00
	Balanced	0.25	0.25	0.50
	LPS	0.00	0.00	1.00

As noted in Section 3.3.4, the metrics are scaled to between 0 and 1 by using practical ranges that are defined by the user. The selected ranges are partially informed by the results of the meta-parameter analysis, but some trial-and-error is required, particularly in the cases with a balanced preference, to achieve the best possible objectives for different case studies.

#### 3.4.4 Computational tests

To evaluate the utility of the proposed approach as objectively as possible, the calibration and validation performance obtained using the approach is compared with that achieved using two benchmark performance models that replicate common urban growth strategies. These models include the Random Constraint Match (RCM) neutral model (RIKS, 2011), which generates reference maps characterised by a speckled distribution of small clusters of each land-use class, and the Growing Clusters (GC) neutral model (Van Vliet et al., 2013b), which generates reference maps characterised by agglomerated distributions of large clusters of each land-use class (generated using the default neighbourhood rule settings in Metronamica). If the performance of the calibrated LUCA model exceeds that of the benchmark models, the LUCA model can be considered to have captured the processes driving land use change correctly (Hagen-Zanker and Lajoie, 2008). For this research, ten reference maps are generated for each case study for both the calibration and validation period, and reference metrics calculated for each reference map.

To evaluate the output 50 model replicates with different LUCA model random seeds are run with the set of neighbourhood weighting parameters obtained at the end of the proposed calibration approach for the calibration (1990-2000) and validation (2000-2006) periods. For

each simulated output map, a number of performance metrics are calculated. First, the average metrics of the simulated output are compared with the average benchmark metrics, to determine whether the benchmarks are outperformed. In this case, this corresponds to the average FK or FKS values for the simulated output being greater than the average benchmark values, and the average AWCE value being less than the average benchmark values. The metric values obtained are then compared with the corresponding metrics for the benchmark models using Welch's t-test of statistical significance (Welch, 1947), a method of comparing two independent samples with varying size and variance that obey parametric assumptions, at the 95% confidence limit, to determine whether the average of the output metrics is significantly better or worse than the benchmark metrics obtained from a statistical perspective.

### **3.5 Results and discussion**

This section presents the results of the application of the computational tests outlined in Section 3.4.4. The results demonstrate how to determine a final calibrated model for each case study based on the different configurations tested. To do this, first an objective evaluation is conducted for the calibration and validation periods in comparison to the benchmark models. The performance is summarised in Figures 3.12 and 3.13 respectively, with green colouring indicating that performance of the models calibrated with the proposed approach is statistically significantly better (at the 95% confidence limit) than that of both benchmark models, yellow that model performance is superior to one benchmark and inferior to the other, and red that model performance is inferior to both benchmarks. In Figures 3.12 and 3.13 the average metric value across the 50 model replicates is shown, as there tended to be limited spread across the different replicates (see box-plots in Supplementary material 3D). Configurations that outperform all benchmarks are then evaluated for their physical plausibility. First, the simulated output maps are compared with the data using visual interpretation. At this stage, if multiple configurations have resulted in superior performance, a solution is selected that produces the simulated output most consistent with process knowledge. The parameters of this solution are then evaluated against discursive knowledge, to determine if the calibrated model is consistent with expectation.

#### **3.5.1 Objective evaluation**

##### **3.5.1.1 Calibration performance**

The performance of the different configurations applied to the case studies for the calibration period is summarised in Figure 3.12. As shown, the performance is generally significantly better than the benchmarks for the calibration period, irrespective of the metrics used and the

preference for a particular objective, reflected by a majority of the cells being green. This highlights the overall robustness of the proposed approach. A detailed discussion of the impact of the choice of metrics and preference on the robustness of the proposed approach is given in the subsequent sections.

Preference		Locational agreement			Balanced			Landscape pattern structure		
Metrics	Case	FK	FKS	AWCE	FK	FKS	AWCE	FK	FKS	AWCE
FK vs AWCE	Berlin	0.947	0.053	0.008	0.945	0.046	0.005	0.938	0.014	0.001
	Budapest	0.923	0.076	0.017	0.921	0.086	0.009	0.917	0.047	0.001
	Lisbon	0.894	0.087	0.019	0.892	0.078	0.006	0.876	0.011	0.002
	Madrid	0.919	0.186	0.007	0.907	0.138	0.002	0.894	0.124	0.001
FKS vs AWCE	Berlin	0.917	0.065	0.035	0.946	0.061	0.003	0.938	0.010	0.001
	Budapest	0.881	0.100	0.040	0.918	0.077	0.006	0.915	0.047	0.001
	Lisbon	0.872	0.128	0.025	0.872	0.102	0.004	0.874	0.011	0.001
	Madrid	0.907	0.198	0.014	0.902	0.139	0.001	0.894	0.124	0.001
FK & FKS vs AWCE	Berlin	0.947	0.068	0.006	0.947	0.062	0.007	0.938	0.010	0.001
	Budapest	0.921	0.094	0.015	0.919	0.082	0.008	0.915	0.047	0.001
	Lisbon	0.886	0.106	0.021	0.891	0.082	0.006	0.876	0.011	0.002
	Madrid	0.905	0.198	0.013	0.913	0.192	0.005	0.894	0.124	0.001

Figure 3.12. Average metric values for 50 model replicates compared to two sets of benchmark model metrics using statistical significance testing at the 95% confidence level for the calibration period. Cells are coloured by performance compared to the benchmark, green indicates the average is significantly superior to both benchmarks, yellow indicates the average is significantly superior to one benchmark, and red indicates the average is inferior to both benchmarks.

The performance across all three metrics (FK, FKS and AWCE) is particularly good when locational agreement and landscape pattern structure are balanced, as there are only two cases where both benchmarks are not exceeded significantly (Figure 3.12). These results suggest that the proposed approach can identify parameter combinations that result in modelled land-use dynamics that successfully balance locational agreement and landscape pattern structure.

When locational agreement is favoured during the fine adjustment stage, values of FK and FKS are generally higher than when landscape pattern structure is favoured or a balanced approach is used, as shown in Figure 3.12. However, this comes at the expense of landscape pattern structure accuracy, the performance of which deteriorates (i.e. values of AWCE are increased) to the point where values are only significantly better than all benchmark values in half the results. This indicates that it is difficult to achieve results that satisfy both locational agreement and landscape pattern structure when only locational agreement is considered during the fine adjustment process and that a balanced approach is needed to achieve acceptable performance

for both locational agreement and landscape pattern structure, which is also consistent with the findings of Newland et al. (2018a).

Values of FKS are significantly better than all benchmarks for all case studies and objective function combinations when locational agreement is favoured during the fine adjustment process. However, somewhat surprisingly, this is not the case when FK is used as the performance metric. As can be seen from Figure 3.12, when FKS and AWCE are used as objectives, FK values are only significantly better than all benchmarks for the Madrid case study, and for Berlin and Budapest are significantly worse than both benchmarks. This likely occurs because the use of FKS as the sole objective places an over-emphasis on capturing the transitions that are present in the data, with insufficient weight given to inertia. This is especially important for cases where there is low or moderate growth, such as Berlin and Budapest. However, in cases with high growth, such as Madrid, this impact is less pronounced.

When there is a preference for landscape pattern structure during the fine adjustment process, values of AWCE are superior (i.e. less) to those obtained when there is a preference for a different objective, as shown in Figure 3.12. However, the singular focus on landscape pattern structure has a detrimental impact on the metrics used to measure locational agreement. This is most noticeable for the FK values, where both benchmarks are only exceeded in three cases, all for the same case study (Budapest).

### *3.5.1.2 Validation performance*

The general performance of the different configurations applied to the different case studies for the validation period is summarised in Figure 3.13. As shown, the results for the validation period generally follow the same trends as for the calibration period in that a balanced preference between locational agreement and landscape pattern structure results in the best overall results, while solely favouring locational agreement or landscape pattern structure results in a deterioration in performance of the other objective, to the point that it is much less likely that corresponding benchmarks are exceeded. Specifically, when a balanced approach is used in conjunction with either FK and AWCE or FK, FKS and AWCE as objectives, validation performance is significantly better than that of all benchmarks for all metrics for three of the four case studies considered, as well as outperforming both benchmarks in most cases for the other case study and only performing significantly worse than both benchmarks in one out of twenty four cases (i.e. FK for Budapest when a balanced approach is used in conjunction with FK and AWCE as objectives). This suggests that the proposed approach is capable of

generating results for both calibration and validation periods that perform significantly better (at the 95% confidence level over 50 replicates) for all three performance metrics in almost all experiments, provided both locational agreement and landscape pattern structure are appropriately balanced during the fine adjustment process, and FKS is not used as the sole objective quantifying locational agreement.

Preference		Locational agreement			Balanced			Landscape pattern structure		
Metrics	Case	FK	FKS	AWCE	FK	FKS	AWCE	FK	FKS	AWCE
FK vs AWCE	Berlin	0.970	0.013	0.009	0.972	0.012	0.005	0.970	0.003	0.002
	Budapest	0.899	0.047	0.018	0.899	0.054	0.019	0.902	0.045	0.013
	Lisbon	0.888	0.052	0.019	0.886	0.042	0.011	0.879	0.011	0.006
	Madrid	0.937	0.093	0.008	0.934	0.080	0.006	0.926	0.064	0.007
FKS vs AWCE	Berlin	0.946	0.022	0.026	0.968	0.006	0.004	0.970	0.008	0.003
	Budapest	0.863	0.064	0.036	0.902	0.050	0.016	0.901	0.040	0.011
	Lisbon	0.873	0.077	0.029	0.874	0.054	0.011	0.875	0.004	0.009
	Madrid	0.914	0.089	0.014	0.929	0.073	0.006	0.926	0.064	0.007
FK & FKS vs AWCE	Berlin	0.969	0.013	0.004	0.970	0.013	0.005	0.970	0.008	0.003
	Budapest	0.898	0.051	0.018	0.907	0.058	0.010	0.901	0.040	0.011
	Lisbon	0.882	0.067	0.024	0.886	0.046	0.013	0.879	0.011	0.006
	Madrid	0.911	0.093	0.014	0.932	0.088	0.008	0.926	0.064	0.007

Figure 3.13. Average metric values for 50 model replicates compared to two sets of benchmark model metrics using statistical significance testing at the 95% confidence level for the validation period. Cells are coloured by performance compared to the benchmark green indicates the average is significantly superior to both benchmarks, yellow indicates the average is significantly superior to one benchmark, and red indicates the average is inferior to both benchmarks.

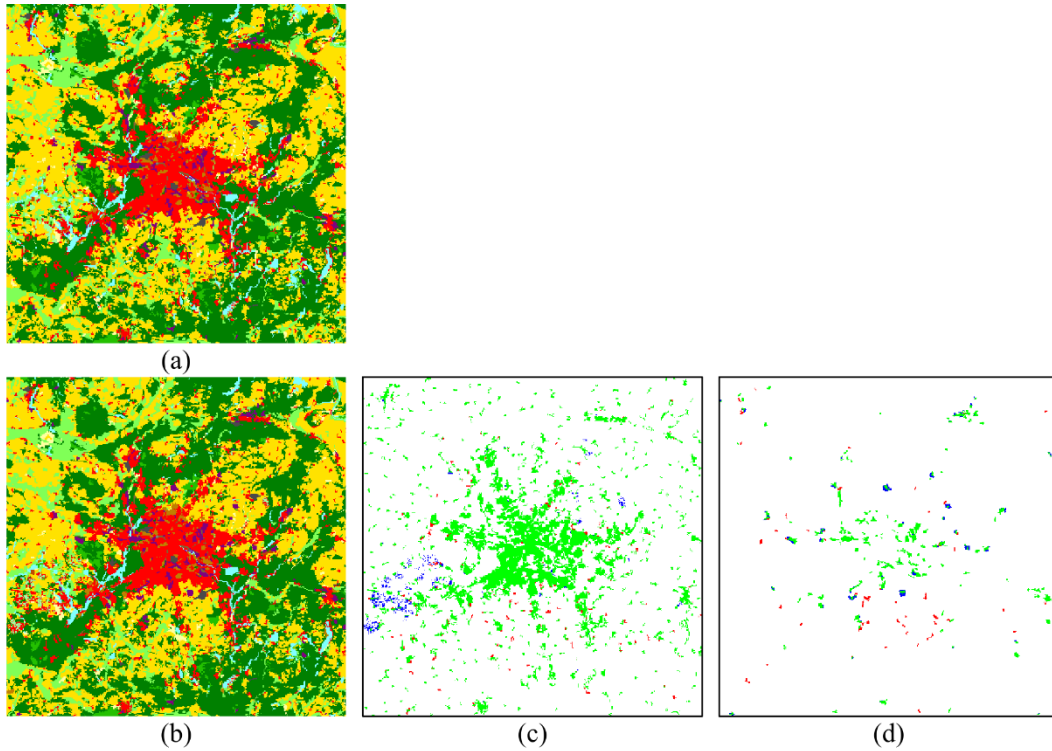
While the validation performance of the models calibrated is significantly better than that of both benchmarks when an appropriate balance of objectives and preferences is used, validation performance is generally worse than calibration performance, as can be expected, given that performance is tuned to the calibration data. This deterioration in performance is particularly pronounced for the case studies experiencing less growth (i.e. Berlin and Budapest) and when FKS is used as the sole objective for locational agreement. This can be explained by the fact that FKS only focuses on transitioned cells, which can result in an over-emphasis on capturing the small number of cells that do in fact transition in low-growth cases during the calibration period, at the expense of capturing inertia.

### 3.5.2 Simulated output evaluation

This section presents an evaluation of a simulated output map for a configuration that performs statistically significantly better than the benchmarks for each case study. The simulated output maps are compared with the data via visual inspection, to determine if the resultant simulated

output is sufficiently realistic, further verifying the quality of the calibrated model. Overall, the results for the four case studies demonstrated that the proposed approach produced realistic output maps provided an appropriate balance between locational agreement and landscape pattern structure is achieved, as shown in Figures 3.14 to 3.17, and discussed below.

Figure 3.14 shows the similarity between the simulated output and the data for the Berlin case study resulting from a balanced preference with FK and AWCE as the objectives. The simulated output appears fairly similar to the data, however, there is some variation between the two maps, with the larger amount of *residential* (red) cells allocated in the western region of the simulated output map, as shown in Figure 3.14c. Also, the simulated map for *industry & commerce* (Figure 3.14d) tends to show slightly larger clusters in the simulated output compared to the data.



### Legend

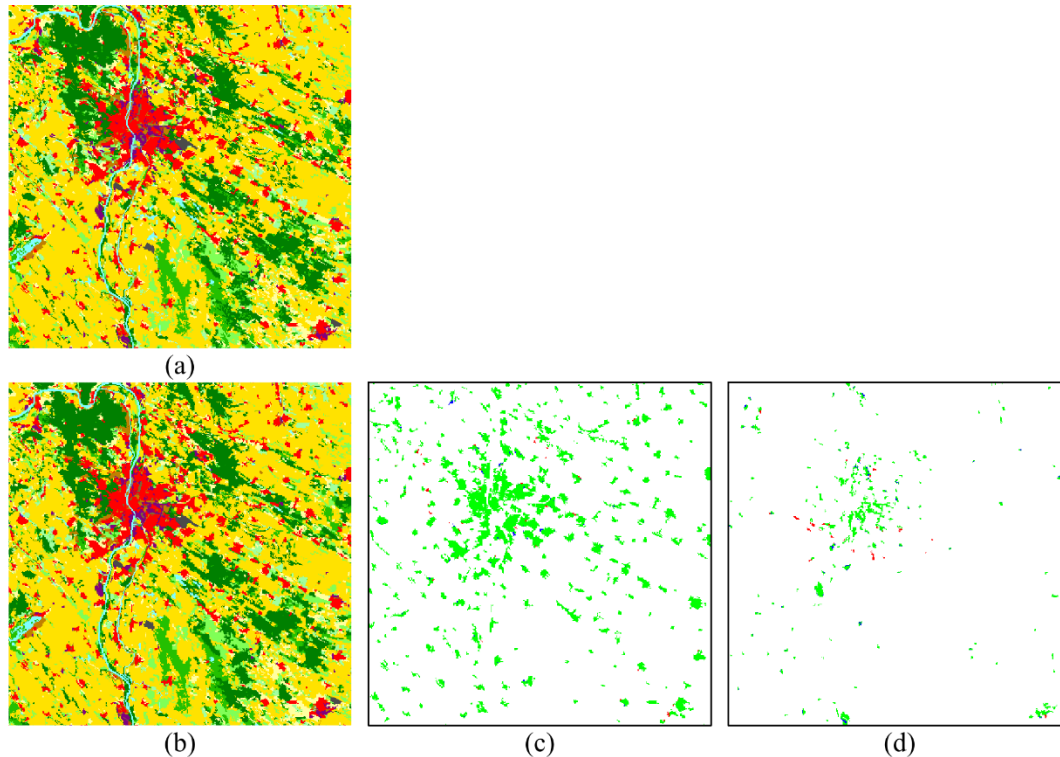
Land-use map			Agreement map				
	Natural areas		Residential		Seaports		In neither map
	Arable land		Industry & commerce		Airports		In both maps
	Permanent crops		Recreation areas		Mine & dump sites		Only in data map
	Pastures		Forest		Fresh water		Only in simulated map
	Other agriculture		Road & rail		Marine water		

Figure 3.14. Berlin data compared with simulated output for the calibration period for best performing objective configuration. (a) is the data map, (b) is the simulated output map, (c) is the agreement of the class residential (red) between maps (a) and (b), and (d) is the agreement of the class industry & commerce (purple) between maps (a) and (b).

The similarity between the simulated output and the corresponding data for the Budapest case study is shown in Figure 3.15 when a balanced preference is used with FK, FKS and AWCE as the objectives. In the Budapest case study, moderate land-use change has taken place over the calibration period, and the results obtained show the model captured the inertia well. For the changes that did occur, the model did not always allocate the transitions to the correct locations, but the size of the simulated clusters resembles the data. This is shown in Figure 3.15, as there is good agreement between the patterns and locations of the urban classes, *residential* (red) and *industry & commerce* (purple), as shown in Figures 3.15c and 3.15d, as well as *recreation areas* (brown). This is also true for the major agricultural class, *arable land* (yellow). There is some variation in the amount of interspersion of the classes *natural areas*



(light green) and *forest* (dark green), with the simulated output producing more clumped areas, though this is quite minor.



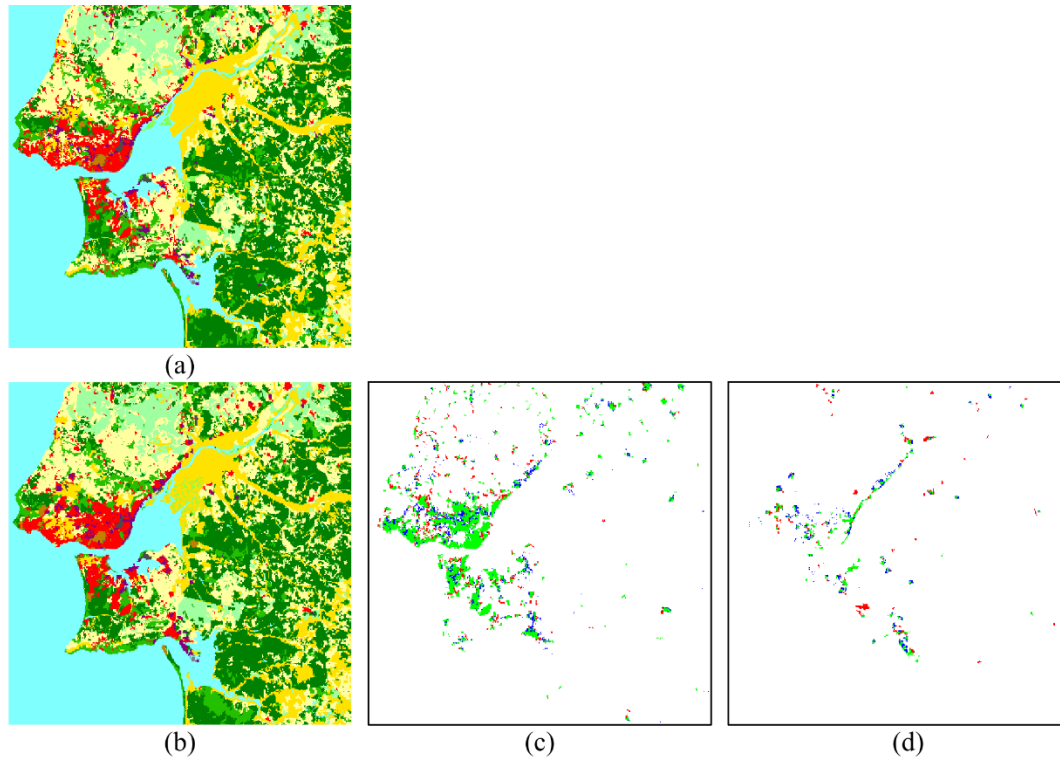
### Legend

Land-use map			Agreement map				
	Natural areas		Residential		Seaports		In neither map
	Arable land		Industry & commerce		Airports		In both maps
	Permanent crops		Recreation areas		Mine & dump sites		Only in data map
	Pastures		Forest		Fresh water		Only in simulated map
	Other agriculture		Road & rail		Marine water		

Figure 3.15. Budapest data compared with simulated output for the calibration period for best performing objective configuration. (a) is the data map, (b) is the simulated output map, (c) is the agreement of the class residential (red) between maps (a) and (b), and (d) is the agreement of the class industry & commerce (purple) between maps (a) and (b).

For the Lisbon case study, the simulated output for a balanced preference with FK, FKS and AWCE as the objectives is shown in Figure 3.16. As shown, the general behaviour of the simulated output is sufficiently consistent with the data, though there are some differences in the newly allocated land-use classes. The simulated *residential* locations tend to fill in space in the urban core, creating larger clusters than in the data. This is shown in Figure 3.16c, as the red cells tend to be at the edge of the existing *residential* area, whereas the blue tends to be within the existing *residential* area. The simulated cluster size of the *industry & commerce* area

resembles the data and while the allocation takes place in similar areas, exact matches were not often found, as shown in Figure 3.16d. Also, the high degree of interspersion of the different agricultural and natural land-use classes has resulted in a noticeable amount of the class *pastures* (light green) interspersed amongst a region of *arable land* (dark yellow) in the centre of the map for the simulated output that is not present in the data.



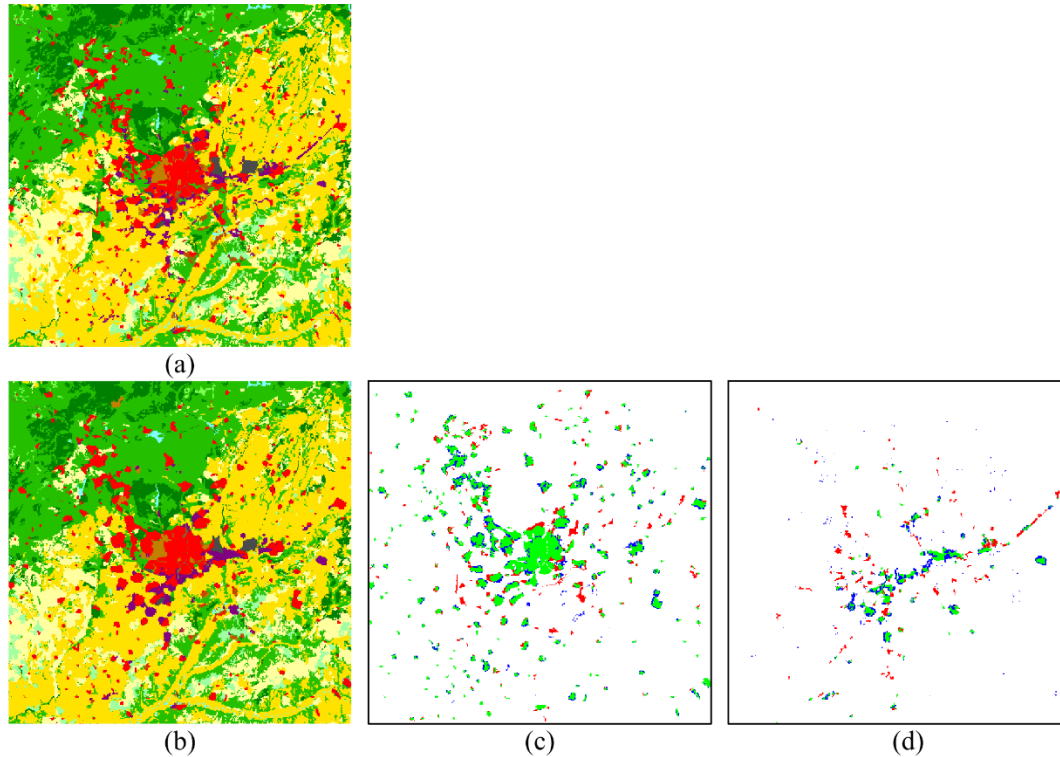
### Legend

Land-use map			Agreement map	
Natural areas	Residential	Seaports		In neither map
Arable land	Industry & commerce	Airports		In both maps
Permanent crops	Recreation areas	Mine & dump sites		Only in data map
Pastures	Forest	Fresh water		Only in simulated map
Other agriculture	Road & rail	Marine water		

Figure 3.16. Lisbon data compared with simulated output for the calibration period for best performing objective configuration. (a) is the data map, (b) is the simulated output map, (c) is the agreement of the class residential (red) between maps (a) and (b), and (d) is the agreement of the class industry & commerce (purple) between maps (a) and (b).

The consistency between the simulated output and corresponding data for the Madrid case study is shown in Figure 3.17 for a balanced preference with FK, FKS and AWCE as the objectives. As shown, there is reasonably good agreement for the expansion of the existing central urban region, composed of *residential* (red), *industry & commerce* (purple) and *recreation areas* (brown). The main variation is that the urban classes in the simulated output

appear slightly more clumped than in the data. As shown in Figures 3.17c and 3.17d, there tends to be a more speckled distribution of red cells, whereas the blue cells tend to be at the fringes of the existing urban region.



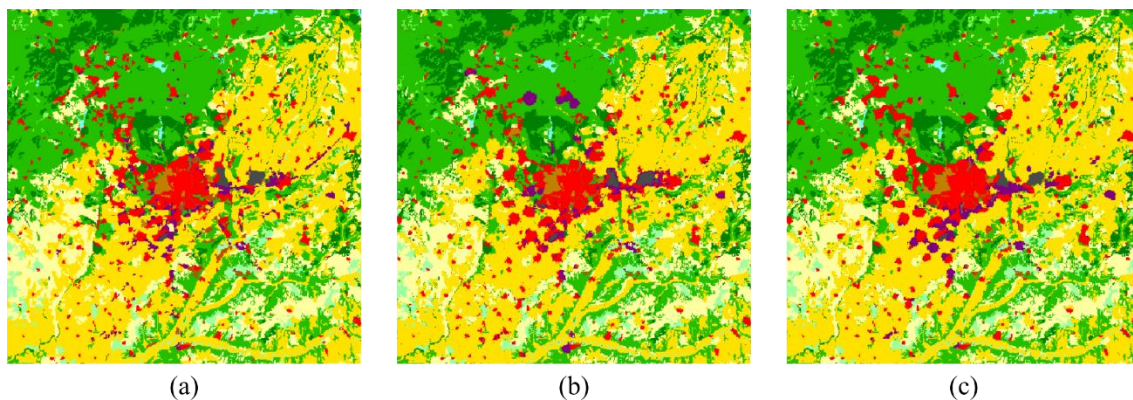
### Legend

Land-use map			Agreement map				
	Natural areas		Residential		Seaports		In neither map
	Arable land		Industry & commerce		Airports		In both maps
	Permanent crops		Recreation areas		Mine & dump sites		Only in data map
	Pastures		Forest		Fresh water		Only in simulated map
	Other agriculture		Road & rail		Marine water		

Figure 3.17. Madrid data compared with simulated output for the calibration period for best performing objective configuration. (a) is the data map, (b) is the simulated output map, (c) is the agreement of the class residential (red) between maps (a) and (b), and (d) is the agreement of the class industry & commerce (purple) between maps (a) and (b).

It is worth highlighting that for the low and moderate growth case studies, Berlin and Budapest, respectively, there was only one configuration that performed statistically significantly better than the benchmarks, meaning only a single solution map required visual inspection and interpretation. However, for the high growth case studies (Lisbon and Madrid), more configurations resulted in superior benchmark performance. For these situations, visual comparison of the simulated output is important to determine the best performing calibrated model. An example of this is presented in Figure 3.18, which shows another configuration for

the Madrid case, using a balanced preference with FK and AWCE as objectives (Figure 3.18b), compared with the data (Figure 3.18a), and the solution that was selected (Figure 3.18c). As shown, there is a large variation in the output obtained. Most notably, the alternate solution (Figure 3.18b) results in a large, clustered formation of the class *industry & commerce* (purple) in the north region of the map that does not appear in the data, and as such is less realistic than the simulated output for the other configuration. This highlights the sensitivity of the output to the configuration used, and illustrates why visual inspection is beneficial as further interpretation of the output that is not captured with objective assessment alone.



*Figure 3.18. Comparison of simulated output for different Madrid configurations that outperformed benchmarks. (a) is the data map, (b) is the simulated output for a balanced preference with FK and AWCE as objectives, and (c) is the simulated output for a balanced preference with FK, FKS and AWCE as objectives*

### 3.5.3 Parameter analysis

This section presents an analysis of the parameters for the four solutions corresponding to the best objective performance and most realistic simulated output. The parameter analysis is used to further verify that the obtained solutions are consistent with process understanding. In general, the parameters obtained were consistent with expectation, because the neighbourhood rules were parameterised to generate shapes that were consistent with process knowledge. However, there were cases where rules were included (as discussed previously in Section 3.4.3.1) that were not necessarily consistent with process knowledge, which could be addressed with additional manual intervention in the parameter elimination stage. Given the variation in the interaction elimination across the case studies, and the large number of parameters that were calibrated, the parameter analysis is focussed on the parameters that were included in each case study, the inertia points and self-influence tails. The final parameter values for each case study are presented in Supplementary Material 3E.



A comparison plot of the inertia point parameters obtained for the analysed solutions is presented in Figure 3.19, ordered by class and case study. As shown, for each case study the classes that exhibit the highest inertia tend to be the urban classes, *residential*, *industry & commerce*, and *recreation areas*, which is consistent with expectation, as such classes have high transition costs. The land-use class *arable land*, which essentially serves as a vacant class to facilitate transitions, exhibits the lowest inertia for each case, which is expected as such land-uses tend to have the lowest transition costs.




Figure 3.19. Comparison of inertia point influence values across classes for the case studies

Figure 3.20 shows a comparison of the self-influence tail parameters (i.e. the influence of class  $i$  in the neighbourhood of class  $i$ ) for the different solutions for each case study. The self-influence tails exhibit behaviour that is also consistent with expectation. The urban classes exhibit a higher degree of self-influence across the different case studies, agricultural (*pastures*, *permanent crops*, *agricultural areas*) classes exhibit less self-influence, and the class *arable land* exhibits virtually no self-influence. A trend observed across the case studies is that the higher growth cases (Lisbon and Madrid) exhibit higher self-influence for the class *residential* than the low and moderate growth cases (Berlin and Budapest). This would be expected, as higher growth cases will have more newly allocated *residential* cells, which will require a stronger attraction to ensure they are allocated near the existing *residential* region. Hence,

based on the parameter evaluation, the calibration method was able to generate parameters that are consistent with expectation.

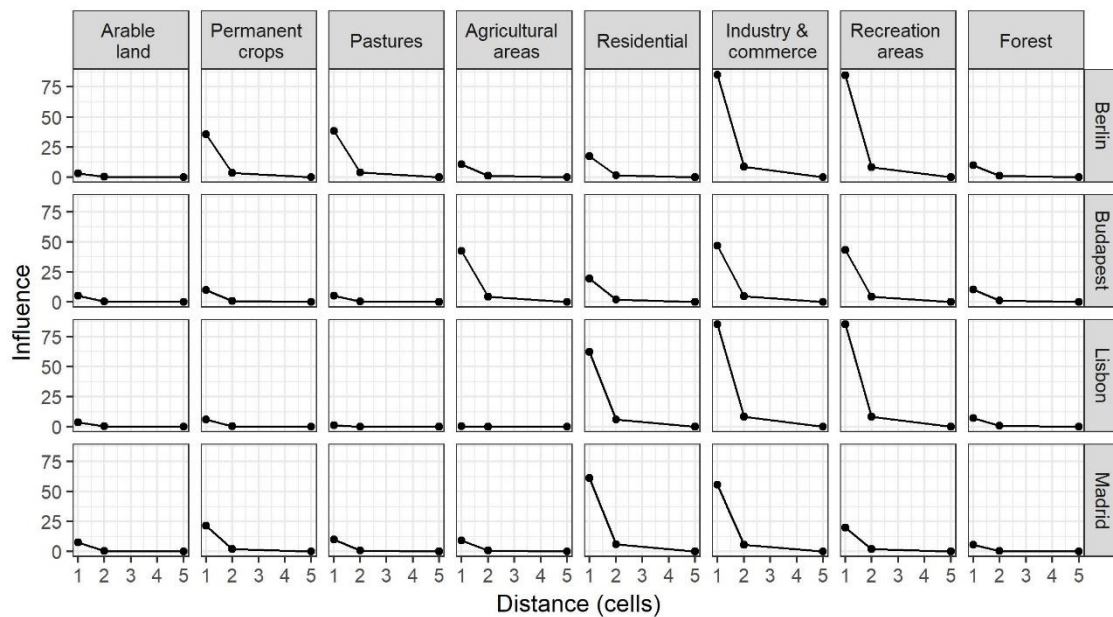


Figure 3.20. Comparison of self-influence tail parameter values across classes for each case study.

### 3.6 Summary and conclusions

Transition potential LUCA models use neighbourhood rules to replicate the spatial dynamics that drive land-use changes in a region. Neighbourhood rules must be effectively calibrated for model application, as these are the main calibration parameters of such models due to their impact on parameter dimensionality. This paper presents a semi-automatic calibration method that integrates objective analysis with discursive input to facilitate efficient calibration of neighbourhood rules within a limited computational budget, achievable using a desktop PC. The method first reduces the complexity of the calibration problem and then calibrates the remaining neighbourhood rules in a computationally efficient manner, based on a set of metrics that quantify the two key aspects of LUCA model performance; locational agreement and landscape pattern structure.

The utility of the proposed approach was demonstrated via application to four European case studies with varying physical characteristics and rates of growth. For each case study, the method was implemented with a focus on a certain objective, either locational agreement, landscape pattern structure, or a balance between the two. Based on a statistical analysis of the simulated output metrics compared with metrics obtained from two benchmark models, and

consideration of the simulated output maps and parameters, optimal performance required using a balanced objective preference. For this research, the best performance for the low growth case study (Berlin) was achieved when using a balanced approach between two objectives, Fuzzy Kappa and Area-Weighted Clumpiness Error. For the moderate (Budapest) and high (Lisbon and Madrid) growth cases, performance was maximised when using a balanced approach between three objectives, Fuzzy Kappa, Fuzzy Kappa Simulation and Area-Weighted Clumpiness Error, suggesting that locational agreement is best quantified by balancing between the agreement of transitions and the agreement of the entire land-use map.

This research demonstrates the efficiency that the integration of process knowledge affords process-specific calibration methods. The results suggest further improvements in efficiency and the quality of the final output could be achieved by utilising additional process knowledge, with the elimination or inclusion of certain neighbourhood rules, and the input of more complex neighbourhood rule shapes. The results also suggest potential improvement could be achieved by adjusting the calibration objectives. As discussed in Section 3.5.2, there tended to be over-clustering of the resultant output for urban regions. Hence, a different aggregation strategy for combining the class level clumpiness errors may improve the results. Also, additional metrics could be used during the automatic calibration procedure that emphasize different pattern aspects, such as the fractal dimension or edge density (McGarigal, 2014), which were only captured by visual inspection for the case study application. The demonstrated application shows that the proposed approach is a step towards efficient LUCA model calibration with the potential to facilitate more widespread use of LUCA models to support scenario and policy analysis.

The efficiency that is achieved with a process-specific method also has the potential to increase the efficiency of more computationally demanding approaches, such as optimisation. This would further reduce the computational demands of such approaches, as has been demonstrated with the use of meta-modelling by Şalap-Ayça et al. (2018). Integrating a process-specific method with optimisation would have the added benefit of integrating further process knowledge into such methods, potentially improving the applicability of the resulting models.

### **3.7 Acknowledgements**

The authors wish to acknowledge the financial support from the Bushfire and Natural Hazards Cooperative Research Centre, made available by the Commonwealth of Australia through the Cooperative Research Centre program.

### 3.8 Supplementary material

#### 3.8.1 Supplementary material 3A: Derivation of empirical measures

This section provides a detailed derivation of the empirical measures used in the interaction elimination and parameter categorisation and initialisation sections of the proposed approach. The link between data from the contingency table and the enrichment factor at distance zero is also shown.

The generation of the contingency table requires an analysis of the state of each cell in two data maps:

$$S_{i,j}(\widehat{\mathbf{X}}_0, \widehat{\mathbf{X}}_a) = \{c \in C: \hat{x}_{c,0} = i, \hat{x}_{c,a} = j\} \quad (3A.1)$$

where  $S_{i,j}$  logs the state of the land-use class in the cell  $c$  of interest at the start and end of the calibration period. Populating the contingency table requires evaluating the state transition for each possible combination of land-use classes for the entire land-use map:

$$\eta_{i,j} = |S_{i,j}| \quad (3A.2)$$

where  $\eta_{i,j}$  is the total numbers number of cells that are class  $i$  in  $\widehat{\mathbf{X}}_0$  and class  $j$  in  $\widehat{\mathbf{X}}_a$ . The case where  $i$  does not equal  $j$  corresponds to a transition from  $i$  to  $j$ , and the case where  $i$  is equal to  $j$  ( $\eta_{i,i}$ ) shows the total number of cells that did not change land-use class between  $\widehat{\mathbf{X}}_0$  and  $\widehat{\mathbf{X}}_a$ . With these values determined the contingency table is populated such that:

Table 3A.1. Example contingency table, populated by logging the land-use class in each cell between two maps

		$\widehat{\mathbf{X}}_a$			
		$A_1$	$A_2$	...	$A_n$
$\widehat{\mathbf{X}}_0$	$A_1$	$\eta_{1,1}$	$\eta_{1,2}$	...	$\eta_{1,n}$
	$A_2$	$\eta_{2,1}$	$\eta_{2,2}$	...	$\eta_{2,n}$
	$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$
	$A_n$	$\eta_{n,1}$	$\eta_{n,2}$	...	$\eta_{n,n}$

From the contingency table, the empirical measures inertia rate and conversion rate can be calculated. The inertia rate, which quantifies the tendency of a particular class to persist, is calculated by:

$$IR_i = \frac{\eta_{i,i}}{\sum_{m=1}^n \eta_{i,m}} \quad (3A.3)$$

The conversion rate, which quantifies the tendency of transitions to a particular class as a function of all transitions to that class, is calculated by:



$$CR_{i,j} = \frac{\eta_{i,j}}{(\sum_{m=1}^n \eta_{m,j}) - \eta_{i,i}} \quad (3A.4)$$

To calculate the enrichment factor requires evaluating the composition of the neighbourhoods of cells that transitioned. The size of the neighbourhood is defined by a radius  $d_{MAX}$ , the maximum allowable distance between two cells for inclusion in the neighbourhood. The neighbourhood is divided into a set of discrete unit-distance rings (i.e. distances of 1, 2, 3), based on the nearest unit-distance between the neighbourhood cell and the central cell. This is shown in Figure 3A.1 for a maximum cellular distance of three cells. Cells are allocated to the nearest unit-distance neighbourhood ring, so the neighbourhood of distance 0 comprises 1 cell at the location of the transition, distance 1 comprises a total of 8 cells at distances of 1 and 1.41 cell lengths from the centre, the neighbourhood of distance 2 comprises a total of 12 cells at distances 2 and 2.24 cell lengths from the centre, and the neighbourhood of distance 3 comprises a total of 8 cells at distances 2.83 and 3 cell lengths from the centre.




Figure 3A.1. Example delineation of a neighbourhood of maximum distance three cell into a set of unit-distance rings for counting neighbourhood composition

Following the subdivision of the neighbourhood into discrete unit-distance rings, the number of cells of each land-use class in each unit-distance ring in  $X_0$  is tallied:

$$n_{i,j,d} = |\{c \in D_d(c), x_{0,c} = j\}| \quad (3A.5)$$

where  $n_{i,j,d}$  is the number of cells of land-use class  $j$  in the neighbourhood of distance  $d$  of a cell that transitioned to land-use class  $i$ , and  $D_d(c)$  is the set of cells in the neighbourhood of cell  $c$  at distance  $d$ . This tally of absolute neighbourhood representation of a land-use class at a certain distance from cells that transitioned to a certain land-use class is then converted to a relative representation, expressing the absolute representation as the percentage of the maximum number of possible cells in the neighbourhood (e.g. for distance one, 1 cell in the

neighbourhood is equivalent to a relative representation of 0.125, 2 cells in the neighbourhood is equivalent to a relative representation of 0.250).

$$R_{i,j,d} = \frac{n_{i,j,d}}{n_d} \quad (3A.6)$$

where  $R_{i,j,d}$  is the relative neighbourhood representation of the neighbourhood of distance  $d$  occupied by land-use  $j$  for cells that transitioned to land-use class  $i$ , and  $n_d$  is the total number of cells in the neighbourhood of distance  $d$ . From this, the average relative representation is computed for each cell that transitioned to a particular class:

$$\bar{R}_{i,j,d} = \frac{1}{|\tilde{C}_i|} \sum_{i \in \tilde{C}_i} R_{i,j,d} \quad (3A.7)$$

where  $\bar{R}_{i,j,d}$  is the average relative representation of land-use class  $j$  in the neighbourhood at distance  $d$  of cells that transitioned to land-use class  $i$ , and  $\tilde{C}_i$  is the set of cells that transitioned to land-use class  $i$ . The enrichment factor is then computed by:

$$EF_{i,j,d} = \log_{10} \left( \frac{\bar{R}_{i,j,d}}{\frac{n_j}{n}} \right) \quad (3A.8)$$

where  $EF_{i,j,d}$  is the enrichment factor for the presence of land-use class  $j$  in the neighbourhood at a distance  $d$  of cells that transitioned to land-use class  $i$ ,  $n_j$  is the number of cells of the neighbourhood class  $j$  in the land-use map  $\hat{X}_0$  and  $n$  is the total number of cells in the land-use map. For ease of interpretation, enrichment factor values are log-scaled by a factor of ten, with negative values indicating under-representation, and positive values indicating over-representation.

It is possible to determine the enrichment factor value at distance 0 by using the contingency table. This is calculated by:

$$EF_{i,j,0} = \frac{\frac{\eta_{j,i}}{\left(\sum_{m=1}^n \eta_{m,i}\right) - \eta_{i,i}}}{\frac{n_j}{N}} \quad (3A.9)$$

Hence, it is possible to derive the enrichment factor value at distance zero, which provides information about conversions, using the contingency table. However, this highlights why the conversion rate can be more effective at capturing the different conversions that are occurring in the data. For example, many conversions may be occurring from a certain class to another, but the enrichment factor may not suggest over-representation, indicating that a conversion

point is not required, because there is a large representation of the class in the landscape. A common example of this are conversions from large, relatively passive land-use classes such as natural vegetation, because they occupy a large area of the landscape, and facilitate many conversions to different classes, which are meaningful to include.

### 3.8.2 Supplementary material 3B: Functional dependency of parameter categorisation scheme

This section details how the proposed meta-parameterisation and categorisation, detailed in Section 3.3.2, control the relative influence of the different interactions types, and the functional dependence of the neighbourhood potential of the normalised weighting parameters within each interaction type.

The proposed approach introduces three meta-parameters ( $\theta_{CP}$ ,  $\theta_{ST}$ ,  $\theta_{CT}$ ) that describe the inter-type importance of each type of interaction  $P'_{CP}$ ,  $P'_{ST}$ , and  $P'_{CT}$  with respect to the inertia point interaction type  $P'_{IP}$ , and a quantisation scheme that describes the intra-type importance within each interaction type. The impact of this parameter representation on the neighbourhood potential is given by:

$$N_{c,i,t} = N_{c,i,t}^{IP} + \theta_{CP}N_{c,i,t}^{CP} + \theta_{ST}N_{c,i,t}^{ST} + \theta_{CT}N_{c,i,t}^{CT} \quad (3B.1)$$

where  $N_{c,i,t}^k$  is the neighbourhood potential of interaction type  $k$  for cell  $c$  to support land-use type  $i$  at time  $t$ . Given the implemented parameterisation, the individual interaction-type terms can be defined as below.

For the inertia points:

$$N_{c,i,t}^{IP} = \begin{cases} \tilde{w}_{i,i,0} & \text{If } X_{c,t} = i \\ 0 & \text{Otherwise} \end{cases} \quad (3B.2)$$

For the conversion points:

$$N_{c,i,t}^{CP} = \begin{cases} \tilde{w}_{i,X_{c,t},0} & \text{If } X_{c,t} \neq i \\ 0 & \text{Otherwise} \end{cases} \quad (3B.3)$$

For the self-influence tails:

$$N_{c,i,t}^{ST} = \sum_{\substack{c' \in D(c) \\ i, X_{c',t} \in P'_{IT}}} \tilde{w}_{i,i,d(c,c')} \quad (3B.4)$$

For the cross-influence tails:

$$N_{c,i,t}^{CT} = \sum_{\substack{c' \in D(c) \\ i, X_{c',t} \in P'_{CT}}} \tilde{w}_{i,X_{c',t},d(c,c')} \quad (3B.5)$$

### 3.8.3 Supplementary material 3C: Contingency tables

This section contains contingency tables for the data maps for all case studies for the calibration (1990-2000) and validation (2000-2006) periods.

Table 3.C1. Berlin contingency table, 1990-2000

		Map 2000																
		LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT
Map 1990	LUC																	
	NAT	2440	0	0	13	0	6	7	0	188	0	0	0	0	21	1	0	2676
	ARL	8	51812	148	1482	27	601	379	127	198	0	0	0	0	31	5	0	54818
	PER	0	1285	327	31	0	28	0	19	7	0	0	0	0	0	0	0	1697
	PAS	13	773	0	12937	32	23	5	3	39	0	0	0	0	0	18	0	13843
	OAG	0	45	0	151	1931	14	8	0	34	0	0	0	0	9	3	0	2195
	RES	4	4	0	3	0	17206	5	15	16	0	0	0	0	2	0	0	17255
	I&C	0	0	0	2	0	0	2033	0	0	0	0	0	0	0	0	0	2035
	REC	0	0	0	0	0	9	0	1572	0	0	0	0	0	0	0	0	1581
	FOR	111	8	0	4	9	58	30	0	58301	0	0	0	1	65	17	0	58604
	R&R	0	0	0	0	0	0	0	0	0	201	0	0	0	0	0	0	201
	POR	0	0	0	0	0	0	0	0	0	0	6	0	0	0	0	0	6
	AIR	12	0	0	0	0	0	0	0	4	0	0	0	559	0	0	0	575
	M&D	20	5	0	0	0	0	0	0	8	0	0	0	0	282	5	0	320
	FRE	5	0	0	0	0	0	0	0	0	0	0	0	0	0	4189	0	4194
	MAR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	TOT	2613	53932	475	14623	1999	17945	2467	1736	58795	201	6	560	410	4238	0	160000	

Table 3.C2. Berlin contingency table, 2000-2006

		Map 2006															
	LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT
Map 2000	NAT	2258	7	0	78	96	0	2	11	150	4	0	0	3	4	0	2613
	ARL	22	52786	40	440	267	127	40	98	58	12	0	5	32	5	0	53932
	PER	0	41	433	0	0	1	0	0	0	0	0	0	0	0	0	475
	PAS	17	270	0	14293	24	4	6	2	5	0	0	0	0	2	0	14623
	OAG	28	2	0	2	1943	0	3	6	0	5	3	0	7	0	0	1999
	RES	15	14	0	34	0	17732	102	25	0	8	1	0	14	0	0	17945
	I&C	19	15	0	7	0	39	2347	6	2	11	0	0	21	0	0	2467
	REC	0	0	0	0	0	19	0	1717	0	0	0	0	0	0	0	1736
	FOR	80	2	0	0	8	31	35	13	58594	8	0	3	20	1	0	58795
	R&R	0	0	0	0	0	21	8	0	0	172	0	0	0	0	0	201
	POR	0	0	0	0	0	0	0	0	0	0	6	0	0	0	0	6
	AIR	52	0	0	30	0	8	4	18	0	0	0	448	0	0	0	560
	M&D	5	1	0	10	12	5	21	9	7	0	0	0	338	2	0	410
	FRE	0	1	0	4	0	0	0	0	6	0	0	0	0	4227	0	4238
	MAR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	TOT	2496	53139	473	14898	2350	17987	2568	1905	58822	220	10	456	435	4241	0	160000

Table 3.C3. Budapest contingency table, 1990-2000

		Map 2000															
	LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT
Map 1990	NAT	7226	31	0	37	0	5	0	0	2522	0	0	0	6	9	0	9836
	ARL	202	77735	329	1256	303	107	102	43	258	61	0	3	85	112	0	80596
	PER	60	356	3460	51	37	6	12	0	13	7	0	0	0	0	0	4002
	PAS	100	465	3	8265	21	25	22	7	17	0	0	0	2	6	0	8933
	OAG	55	36	6	73	9522	22	6	0	36	0	0	0	2	2	0	9760
	RES	6	4	0	16	0	13547	16	11	1	3	0	0	3	7	0	13614
	I&C	1	0	0	0	0	0	2213	1	0	1	0	0	1	0	0	2217
	REC	4	0	0	0	0	3	3	2242	2	0	0	0	0	0	0	2254
	FOR	1503	10	0	15	15	0	0	0	23426	5	0	0	4	5	0	24983
	R&R	0	0	0	0	0	0	0	0	0	271	0	0	0	0	0	271
	POR	0	0	0	0	0	0	0	0	0	0	58	0	0	0	0	58
	AIR	0	0	0	0	0	0	0	0	0	0	0	415	0	0	0	415
	M&D	6	0	0	14	1	0	0	0	0	0	0	0	189	0	0	210
	FRE	0	0	0	2	4	0	0	3	0	0	0	0	0	2842	0	2851
	MAR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	TOT	9163	78637	3798	9729	9903	13715	2374	2307	26275	348	58	418	292	2983	0	160000

Table 3.C4. Budapest contingency table, 2000-2006

		Map 2006															
	LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT
Map 2000	NAT	7001	134	20	316	129	20	3	1	1488	0	0	0	10	41	0	9163
	ARL	351	76176	598	318	223	195	105	19	425	13	0	1	42	171	0	78637
	PER	11	220	3435	14	86	12	2	3	15	0	0	0	0	0	0	3798
	PAS	228	335	32	8825	86	51	54	35	61	0	4	0	8	10	0	9729
	OAG	96	1546	32	308	7221	362	80	81	166	0	0	0	1	10	0	9903
	RES	0	15	0	5	31	13609	44	3	3	3	0	0	0	2	0	13715
	I&C	0	23	0	15	0	25	2294	13	2	2	0	0	0	0	0	2374
	REC	1	5	0	0	0	89	8	2194	1	0	0	0	0	9	0	2307
	FOR	914	36	13	4	34	3	8	5	25251	0	0	0	0	7	0	26275
	R&R	0	1	0	0	3	0	12	0	0	332	0	0	0	0	0	348
	POR	0	0	0	0	0	0	7	0	0	0	51	0	0	0	0	58
	AIR	0	2	0	0	0	0	0	0	0	0	0	416	0	0	0	418
	M&D	0	9	0	13	0	9	2	6	0	0	0	0	236	17	0	292
	FRE	24	2	0	0	11	0	1	2	2	0	0	0	3	2938	0	2983
	MAR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	TOT	8626	78504	4130	9818	7824	14375	2620	2362	27414	350	55	417	300	3205	0	160000



Table 3.C5. Lisbon contingency table, 1990-2000

		Map 2000															
	LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT
Map 1990	NAT	5884	111	43	1	99	232	97	83	2591	31	0	0	74	7	1	9254
	ARL	236	15304	88	186	154	259	105	0	263	9	0	7	19	11	0	16641
	PER	5	312	9558	0	32	34	20	8	17	0	0	0	8	0	0	9994
	PAS	14	1649	0	392	6	2	10	0	2	0	0	0	0	4	2	2081
	OAG	196	366	262	0	30873	1094	241	29	271	44	0	0	43	0	0	33419
	RES	0	0	0	0	0	5372	27	0	0	7	0	2	2	0	0	5410
	I&C	0	0	0	0	0	15	870	11	0	5	0	0	0	0	0	901
	REC	0	0	0	0	0	0	4	397	0	4	0	0	0	0	0	405
	FOR	3794	563	70	0	335	132	120	69	32769	24	0	0	65	5	5	37951
	R&R	0	0	0	0	0	0	0	0	0	21	0	0	0	0	0	21
	POR	0	0	0	0	0	0	0	8	0	0	103	0	0	0	0	111
	AIR	0	0	0	0	0	0	0	0	0	0	0	223	0	0	0	223
	M&D	20	0	4	0	1	2	4	0	2	0	0	0	251	0	0	284
	FRE	53	1	0	0	0	0	0	0	0	0	0	0	0	829	0	883
	MAR	2	0	0	0	0	2	16	4	0	0	10	0	0	0	42388	42422
	TOT	10204	18306	10025	579	31500	7144	1514	609	35915	145	113	232	462	856	42396	160000

Table 3.C6. Lisbon contingency table, 2000-2006

		Map 2006															
	LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT
Map 2000	NAT	6900	119	10	20	335	194	116	73	2327	30	0	0	53	25	2	10204
	ARL	212	16604	20	775	380	167	42	0	91	11	0	0	3	1	0	18306
	PER	33	111	9272	7	385	179	18	1	16	0	0	0	0	3	0	10025
	PAS	4	197	0	352	25	0	0	0	0	0	0	0	0	1	0	579
	OAG	374	216	71	6	28463	1560	97	39	629	23	0	0	17	0	5	31500
	RES	11	2	7	0	33	7006	42	25	7	8	0	0	2	0	1	7144
	I&C	19	0	4	1	17	17	1435	12	0	1	1	0	5	0	2	1514
	REC	0	0	1	0	1	11	8	579	0	0	0	9	0	0	0	609
	FOR	3660	83	18	0	256	87	33	72	31659	26	0	0	18	2	1	35915
	R&R	0	0	0	0	1	1	0	0	0	143	0	0	0	0	0	145
	POR	0	0	0	0	0	0	1	0	0	0	112	0	0	0	0	113
	AIR	0	0	0	0	0	2	0	0	0	0	0	230	0	0	0	232
	M&D	12	11	0	0	13	32	14	0	6	2	0	0	372	0	0	462
	FRE	0	4	0	0	0	0	0	0	0	0	0	0	0	852	0	856
	MAR	1	15	0	0	1	2	1	5	0	0	0	0	0	0	42371	42396
	TOT	11226	17362	9403	1161	29910	9258	1807	806	34735	244	113	239	470	884	42382	160000

Table 3.C7. Madrid contingency table, 1990-2000

		Map 2000																
	LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT	
Map 1990	NAT	46821	268	3	0	50	1051	316	164	357	31	0	12	161	20	0	49254	
	ARL	1540	55265	30	0	383	1492	1059	121	32	70	0	92	351	26	0	60461	
	PER	50	18	4278	0	0	29	3	0	0	0	0	0	18	5	0	4401	
	PAS	13	0	0	624	0	23	5	0	0	0	0	0	0	0	0	665	
	OAG	191	34	15	0	20761	207	45	22	6	5	0	0	101	0	0	21387	
	RES	1	4	0	0	6	7653	85	31	0	1	0	0	0	0	0	7781	
	I&C	3	8	0	0	0	15	1152	5	0	0	0	0	0	0	0	1183	
	REC	0	0	0	0	0	37	0	656	0	0	0	0	0	0	0	693	
	FOR	55	16	0	0	8	21	0	11	12164	0	0	0	0	0	7	0	12282
	R&R	0	0	0	0	0	0	0	0	0	145	0	0	0	0	0	0	145
	POR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	AIR	0	0	0	0	0	0	4	0	0	0	0	0	494	0	0	0	498
	M&D	33	6	0	0	0	29	26	39	0	0	0	0	0	368	2	0	503
	FRE	24	4	0	0	0	0	0	0	0	0	0	0	0	0	719	0	747
	MAR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	TOT	48731	55623	4326	624	21208	10557	2695	1049	12559	252	0	0	598	999	779	0	160000

Table 3.C8. Madrid contingency table, 2000-2006

		Map 2006															
	LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT
Map 2000	NAT	47105	386	17	68	152	451	90	34	166	107	0	6	68	81	0	48731
	ARL	191	52871	35	0	249	865	442	50	5	518	0	186	202	9	0	55623
	PER	17	25	4149	0	45	38	5	0	2	11	0	0	10	24	0	4326
	PAS	13	0	0	586	0	17	0	0	2	4	0	0	0	2	0	624
	OAG	27	58	39	336	20499	109	43	11	17	30	0	0	39	0	0	21208
	RES	116	55	0	0	55	9931	136	86	21	62	0	90	2	3	0	10557
	I&C	16	23	0	0	0	287	2300	30	0	28	0	11	0	0	0	2695
	REC	19	4	0	0	4	47	13	953	1	1	0	0	7	0	0	1049
	FOR	143	4	0	0	16	37	3	6	12324	9	0	15	2	0	0	12559
	R&R	0	0	0	0	0	148	9	0	0	95	0	0	0	0	0	252
	POR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	AIR	6	44	0	0	0	6	4	0	0	0	0	538	0	0	0	598
	M&D	88	78	0	11	25	71	71	41	0	11	0	13	590	0	0	999
	FRE	0	1	0	4	0	0	0	0	0	0	0	0	10	764	0	779
	MAR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	TOT	47741	53549	4240	1005	21045	12007	3116	1211	12538	876	0	859	930	883	0	160000

### 3.8.4 Supplementary material 3D: Boxplots of simulated replicates

This section contains boxplots to show the spread of the results for the different configurations of the proposed approach applied to the different case studies.




Figure 3.D1. Boxplot of simulated replicates for different configurations for Berlin case study for calibration and validation periods




Figure 3.D2. Boxplot of simulated replicates for different configurations for Budapest case study for calibration and validation periods




Figure 3.D3. Boxplot of simulated replicates for different configurations for Lisbon case study for calibration and validation periods




Figure 3.D4. Boxplot of simulated replicates for different configurations for Madrid case study for calibration and validation periods



### 3.8.5 Supplementary material 3E: Final model parameterisations

This section contains the final recommend values for the neighbourhood rules for each case study obtained at the conclusion of the application of the calibration method.

Table 3.E1. Final neighbourhood rule parameter values, Berlin

From	To	Distance			
		0	1	2	5
		Influence			
Natural areas	Arable land	0.00	0.00	0.00	0.00
Natural areas	Permanent crops	0.00	0.00	0.00	0.00
Natural areas	Pastures	0.00	0.00	0.00	0.00
Natural areas	Agricultural areas	0.00	0.00	0.00	0.00
Natural areas	Residential	0.00	0.00	0.00	0.00
Natural areas	Industry & commerce	0.00	0.00	0.00	0.00
Natural areas	Recreation areas	0.00	0.00	0.00	0.00
Natural areas	Forest	7.20	5.57	0.56	0.00
Arable land	Arable land	404.99	2.94	0.29	0.00
Arable land	Permanent crops	51.97	2.94	0.29	0.00
Arable land	Pastures	45.08	1.32	0.13	0.00
Arable land	Agricultural areas	66.56	6.39	0.64	0.00
Arable land	Residential	26.55	1.50	0.15	0.00
Arable land	Industry & commerce	62.11	9.33	0.93	0.00
Arable land	Recreation areas	69.50	1.82	0.18	0.00
Arable land	Forest	25.00	0.00	0.00	0.00
Permanent crops	Arable land	70.32	0.81	0.08	0.00
Permanent crops	Permanent crops	317.63	35.57	3.56	0.00
Permanent crops	Pastures	0.00	15.91	1.59	0.00
Permanent crops	Agricultural areas	0.00	0.00	0.00	0.00
Permanent crops	Residential	94.43	23.92	2.39	0.00
Permanent crops	Industry & commerce	0.00	0.00	0.00	0.00
Permanent crops	Recreation areas	25.23	3.00	0.30	0.00
Permanent crops	Forest	0.00	0.00	0.00	0.00
Pastures	Arable land	25.00	1.32	0.13	0.00
Pastures	Permanent crops	0.00	0.00	0.00	0.00
Pastures	Pastures	500.00	38.20	3.82	0.00
Pastures	Agricultural areas	25.00	3.00	0.30	0.00
Pastures	Residential	12.50	0.00	0.00	0.00
Pastures	Industry & commerce	0.00	0.00	0.00	0.00
Pastures	Recreation areas	0.00	0.00	0.00	0.00
Pastures	Forest	99.50	0.00	0.00	0.00
Agricultural areas	Arable land	0.00	6.39	0.64	0.00
Agricultural areas	Permanent crops	0.00	1.50	0.15	0.00
Agricultural areas	Pastures	11.65	3.00	0.30	0.00
Agricultural areas	Agricultural areas	443.02	10.64	1.06	0.00
Agricultural areas	Residential	0.00	16.22	1.62	0.00
Agricultural areas	Industry & commerce	0.00	1.50	0.15	0.00
Agricultural areas	Recreation areas	0.00	0.00	0.00	0.00
Agricultural areas	Forest	99.50	4.26	0.43	0.00
Residential	Arable land	0.00	0.00	0.00	0.00
Residential	Permanent crops	0.00	0.00	0.00	0.00
Residential	Pastures	0.00	0.00	0.00	0.00
Residential	Agricultural areas	0.00	0.00	0.00	0.00

Residential	Residential	1000.00	17.22	1.72	0.00
Residential	Industry & commerce	0.00	0.00	0.00	0.00
Residential	Recreation areas	12.50	0.00	0.00	0.00
Residential	Forest	12.50	0.00	0.00	0.00
Industry & commerce	Arable land	0.00	0.00	0.00	0.00
Industry & commerce	Permanent crops	0.00	0.00	0.00	0.00
Industry & commerce	Pastures	0.00	0.00	0.00	0.00
Industry & commerce	Agricultural areas	0.00	0.00	0.00	0.00
Industry & commerce	Residential	0.00	0.00	0.00	0.00
Industry & commerce	Industry & commerce	1000.00	84.91	8.49	0.00
Industry & commerce	Recreation areas	0.00	0.00	0.00	0.00
Industry & commerce	Forest	0.00	0.00	0.00	0.00
Recreation areas	Arable land	0.00	0.00	0.00	0.00
Recreation areas	Permanent crops	0.00	0.00	0.00	0.00
Recreation areas	Pastures	0.00	0.00	0.00	0.00
Recreation areas	Agricultural areas	0.00	0.00	0.00	0.00
Recreation areas	Residential	0.00	38.20	3.82	0.00
Recreation areas	Industry & commerce	0.00	0.00	0.00	0.00
Recreation areas	Recreation areas	1000.00	84.60	8.46	0.00
Recreation areas	Forest	0.00	0.00	0.00	0.00
Forest	Arable land	0.00	0.00	0.00	0.00
Forest	Permanent crops	0.00	0.00	0.00	0.00
Forest	Pastures	0.00	0.00	0.00	0.00
Forest	Agricultural areas	25.00	0.00	0.00	0.00
Forest	Residential	12.50	0.00	0.00	0.00
Forest	Industry & commerce	12.50	0.00	0.00	0.00
Forest	Recreation areas	0.00	0.00	0.00	0.00
Forest	Forest	996.23	9.83	0.98	0.00
Road & rail	Arable land	0.00	0.00	0.00	0.00
Road & rail	Permanent crops	0.00	0.00	0.00	0.00
Road & rail	Pastures	0.00	0.00	0.00	0.00
Road & rail	Agricultural areas	0.00	0.00	0.00	0.00
Road & rail	Residential	0.00	0.00	0.00	0.00
Road & rail	Industry & commerce	0.00	0.00	0.00	0.00
Road & rail	Recreation areas	0.00	0.00	0.00	0.00
Road & rail	Forest	0.00	0.00	0.00	0.00
Port area	Arable land	0.00	0.00	0.00	0.00
Port area	Permanent crops	0.00	0.00	0.00	0.00
Port area	Pastures	0.00	0.00	0.00	0.00
Port area	Agricultural areas	0.00	0.00	0.00	0.00
Port area	Residential	0.00	0.00	0.00	0.00
Port area	Industry & commerce	0.00	0.00	0.00	0.00
Port area	Recreation areas	0.00	0.00	0.00	0.00
Port area	Forest	0.00	0.00	0.00	0.00
Airports	Arable land	0.00	0.00	0.00	0.00
Airports	Permanent crops	0.00	27.86	2.79	0.00
Airports	Pastures	0.00	0.00	0.00	0.00
Airports	Agricultural areas	0.00	0.00	0.00	0.00
Airports	Residential	0.00	0.00	0.00	0.00
Airports	Industry & commerce	0.00	25.74	2.57	0.00
Airports	Recreation areas	0.00	0.00	0.00	0.00
Airports	Forest	0.00	0.00	0.00	0.00
Mine & dump sites	Arable land	0.00	0.00	0.00	0.00
Mine & dump sites	Permanent crops	0.00	0.00	0.00	0.00
Mine & dump sites	Pastures	0.00	0.00	0.00	0.00

Mine & dump sites	Agricultural areas	0.00	0.00	0.00	0.00
Mine & dump sites	Residential	0.00	0.00	0.00	0.00
Mine & dump sites	Industry & commerce	0.00	0.00	0.00	0.00
Mine & dump sites	Recreation areas	0.00	0.00	0.00	0.00
Mine & dump sites	Forest	0.00	33.44	3.34	0.00
Fresh water	Arable land	0.00	0.00	0.00	0.00
Fresh water	Permanent crops	0.00	0.00	0.00	0.00
Fresh water	Pastures	0.00	0.00	0.00	0.00
Fresh water	Agricultural areas	0.00	9.83	0.98	0.00
Fresh water	Residential	0.00	0.00	0.00	0.00
Fresh water	Industry & commerce	0.00	0.00	0.00	0.00
Fresh water	Recreation areas	0.00	0.00	0.00	0.00
Fresh water	Forest	0.00	0.00	0.00	0.00
Marine water	Arable land	0.00	0.00	0.00	0.00
Marine water	Permanent crops	0.00	0.00	0.00	0.00
Marine water	Pastures	0.00	0.00	0.00	0.00
Marine water	Agricultural areas	0.00	0.00	0.00	0.00
Marine water	Residential	0.00	0.00	0.00	0.00
Marine water	Industry & commerce	0.00	0.00	0.00	0.00
Marine water	Recreation areas	0.00	0.00	0.00	0.00
Marine water	Forest	0.00	0.00	0.00	0.00

Table 3.E2. Final neighbourhood rule parameter values, Budapest

From	To	Distance			
		0	1	2	5
Natural areas	Arable land	8.13	0.00	0.00	0.00
Natural areas	Permanent crops	0.00	0.00	0.00	0.00
Natural areas	Pastures	8.13	0.00	0.00	0.00
Natural areas	Agricultural areas	0.00	0.00	0.00	0.00
Natural areas	Residential	8.13	0.00	0.00	0.00
Natural areas	Industry & commerce	0.00	0.00	0.00	0.00
Natural areas	Recreation areas	0.00	0.00	0.00	0.00
Natural areas	Forest	32.50	3.00	0.30	0.00
Arable land	Arable land	250.00	5.00	0.50	0.00
Arable land	Permanent crops	32.50	1.50	0.15	0.00
Arable land	Pastures	32.50	1.50	0.15	0.00
Arable land	Agricultural areas	32.50	0.00	0.00	0.00
Arable land	Residential	32.50	0.00	0.00	0.00
Arable land	Industry & commerce	32.50	0.00	0.00	0.00
Arable land	Recreation areas	32.50	0.00	0.00	0.00
Arable land	Forest	8.13	0.00	0.00	0.00
Permanent crops	Arable land	16.25	6.00	0.60	0.00
Permanent crops	Permanent crops	800.89	10.00	1.00	0.00
Permanent crops	Pastures	8.13	1.50	0.15	0.00
Permanent crops	Agricultural areas	8.13	1.50	0.15	0.00
Permanent crops	Residential	8.13	0.00	0.00	0.00
Permanent crops	Industry & commerce	8.13	1.50	0.15	0.00
Permanent crops	Recreation areas	0.00	0.00	0.00	0.00
Permanent crops	Forest	0.00	0.00	0.00	0.00
Pastures	Arable land	32.50	3.00	0.30	0.00
Pastures	Permanent crops	0.00	0.00	0.00	0.00
Pastures	Pastures	500.00	5.00	0.50	0.00
Pastures	Agricultural areas	8.13	0.00	0.00	0.00
Pastures	Residential	16.25	0.00	0.00	0.00
Pastures	Industry & commerce	16.25	0.00	0.00	0.00
Pastures	Recreation areas	16.25	0.00	0.00	0.00
Pastures	Forest	0.00	0.00	0.00	0.00
Agricultural areas	Arable land	8.13	0.00	0.00	0.00
Agricultural areas	Permanent crops	0.00	1.50	0.15	0.00
Agricultural areas	Pastures	8.13	0.00	0.00	0.00
Agricultural areas	Agricultural areas	1000.00	42.45	4.25	0.00
Agricultural areas	Residential	16.25	1.50	0.15	0.00
Agricultural areas	Industry & commerce	8.13	0.00	0.00	0.00
Agricultural areas	Recreation areas	0.00	0.00	0.00	0.00
Agricultural areas	Forest	0.00	0.00	0.00	0.00
Residential	Arable land	0.00	0.00	0.00	0.00
Residential	Permanent crops	0.00	0.00	0.00	0.00
Residential	Pastures	0.00	0.00	0.00	0.00
Residential	Agricultural areas	0.00	0.00	0.00	0.00
Residential	Residential	1000.00	19.35	1.93	0.00
Residential	Industry & commerce	8.13	1.50	0.15	0.00
Residential	Recreation areas	16.25	1.50	0.15	0.00
Residential	Forest	0.00	0.00	0.00	0.00
Industry & commerce	Arable land	0.00	0.00	0.00	0.00
Industry & commerce	Permanent crops	0.00	0.00	0.00	0.00

Industry & commerce	Pastures	0.00	0.00	0.00	0.00
Industry & commerce	Agricultural areas	0.00	0.00	0.00	0.00
Industry & commerce	Residential	0.00	0.00	0.00	0.00
Industry & commerce	Industry & commerce	1000.00	46.71	4.67	0.00
Industry & commerce	Recreation areas	0.00	1.50	0.15	0.00
Industry & commerce	Forest	0.00	0.00	0.00	0.00
Recreation areas	Arable land	0.00	0.00	0.00	0.00
Recreation areas	Permanent crops	0.00	0.00	0.00	0.00
Recreation areas	Pastures	0.00	1.50	0.15	0.00
Recreation areas	Agricultural areas	0.00	0.00	0.00	0.00
Recreation areas	Residential	0.00	1.50	0.15	0.00
Recreation areas	Industry & commerce	0.00	1.50	0.15	0.00
Recreation areas	Recreation areas	1000.00	43.27	4.33	0.00
Recreation areas	Forest	0.00	0.00	0.00	0.00
Forest	Arable land	0.00	0.00	0.00	0.00
Forest	Permanent crops	0.00	0.00	0.00	0.00
Forest	Pastures	0.00	0.00	0.00	0.00
Forest	Agricultural areas	8.13	0.00	0.00	0.00
Forest	Residential	0.00	0.00	0.00	0.00
Forest	Industry & commerce	0.00	0.00	0.00	0.00
Forest	Recreation areas	0.00	0.00	0.00	0.00
Forest	Forest	500.00	10.33	1.03	0.00
Road & rail	Arable land	0.00	0.00	0.00	0.00
Road & rail	Permanent crops	0.00	0.00	0.00	0.00
Road & rail	Pastures	0.00	0.00	0.00	0.00
Road & rail	Agricultural areas	0.00	0.00	0.00	0.00
Road & rail	Residential	0.00	0.00	0.00	0.00
Road & rail	Industry & commerce	0.00	0.00	0.00	0.00
Road & rail	Recreation areas	0.00	0.00	0.00	0.00
Road & rail	Forest	0.00	0.00	0.00	0.00
Port area	Arable land	0.00	0.00	0.00	0.00
Port area	Permanent crops	0.00	0.00	0.00	0.00
Port area	Pastures	0.00	0.00	0.00	0.00
Port area	Agricultural areas	0.00	0.00	0.00	0.00
Port area	Residential	0.00	0.00	0.00	0.00
Port area	Industry & commerce	0.00	0.00	0.00	0.00
Port area	Recreation areas	0.00	0.00	0.00	0.00
Port area	Forest	0.00	0.00	0.00	0.00
Airports	Arable land	0.00	0.00	0.00	0.00
Airports	Permanent crops	0.00	0.00	0.00	0.00
Airports	Pastures	0.00	0.00	0.00	0.00
Airports	Agricultural areas	0.00	0.00	0.00	0.00
Airports	Residential	0.00	0.00	0.00	0.00
Airports	Industry & commerce	0.00	0.00	0.00	0.00
Airports	Recreation areas	0.00	0.00	0.00	0.00
Airports	Forest	0.00	0.00	0.00	0.00
Mine & dump sites	Arable land	0.00	0.00	0.00	0.00
Mine & dump sites	Permanent crops	0.00	1.50	0.15	0.00
Mine & dump sites	Pastures	0.00	3.00	0.30	0.00
Mine & dump sites	Agricultural areas	0.00	0.00	0.00	0.00
Mine & dump sites	Residential	0.00	0.00	0.00	0.00
Mine & dump sites	Industry & commerce	0.00	0.00	0.00	0.00
Mine & dump sites	Recreation areas	0.00	0.00	0.00	0.00
Mine & dump sites	Forest	0.00	0.00	0.00	0.00
Fresh water	Arable land	0.00	0.00	0.00	0.00

Fresh water	Permanent crops	0.00	0.00	0.00	0.00
Fresh water	Pastures	0.00	0.00	0.00	0.00
Fresh water	Agricultural areas	0.00	1.50	0.15	0.00
Fresh water	Residential	0.00	0.00	0.00	0.00
Fresh water	Industry & commerce	0.00	0.00	0.00	0.00
Fresh water	Recreation areas	8.13	3.00	0.30	0.00
Fresh water	Forest	0.00	0.00	0.00	0.00
Marine water	Arable land	0.00	0.00	0.00	0.00
Marine water	Permanent crops	0.00	0.00	0.00	0.00
Marine water	Pastures	0.00	0.00	0.00	0.00
Marine water	Agricultural areas	0.00	0.00	0.00	0.00
Marine water	Residential	0.00	0.00	0.00	0.00
Marine water	Industry & commerce	0.00	0.00	0.00	0.00
Marine water	Recreation areas	0.00	0.00	0.00	0.00
Marine water	Forest	0.00	0.00	0.00	0.00

Table 3.E3. Final neighbourhood rule parameter values, Lisbon

From	To	Distance			
		0	1	2	5
Natural areas	Arable land	8.51	0.00	0.00	0.00
Natural areas	Permanent crops	10.00	1.00	0.10	0.00
Natural areas	Pastures	0.00	0.00	0.00	0.00
Natural areas	Agricultural areas	20.00	1.00	0.10	0.00
Natural areas	Residential	28.68	9.33	0.93	0.00
Natural areas	Industry & commerce	18.34	1.00	0.10	0.00
Natural areas	Recreation areas	20.00	2.00	0.20	0.00
Natural areas	Forest	39.51	4.00	0.40	0.00
Arable land	Arable land	500.00	3.75	0.38	0.00
Arable land	Permanent crops	20.00	1.00	0.10	0.00
Arable land	Pastures	40.00	2.00	0.20	0.00
Arable land	Agricultural areas	20.00	1.00	0.10	0.00
Arable land	Residential	36.38	0.00	0.00	0.00
Arable land	Industry & commerce	26.55	1.63	0.16	0.00
Arable land	Recreation areas	0.00	0.00	0.00	0.00
Arable land	Forest	10.00	0.00	0.00	0.00
Permanent crops	Arable land	20.00	1.00	0.10	0.00
Permanent crops	Permanent crops	1000.00	6.25	0.63	0.00
Permanent crops	Pastures	0.00	0.00	0.00	0.00
Permanent crops	Agricultural areas	10.00	0.00	0.00	0.00
Permanent crops	Residential	0.00	0.00	0.00	0.00
Permanent crops	Industry & commerce	99.50	0.00	0.00	0.00
Permanent crops	Recreation areas	10.00	0.00	0.00	0.00
Permanent crops	Forest	0.00	0.00	0.00	0.00
Pastures	Arable land	38.20	4.00	0.40	0.00
Pastures	Permanent crops	0.00	0.00	0.00	0.00
Pastures	Pastures	851.11	1.32	0.13	0.00
Pastures	Agricultural areas	0.00	0.00	0.00	0.00
Pastures	Residential	0.00	0.00	0.00	0.00
Pastures	Industry & commerce	0.00	0.00	0.00	0.00
Pastures	Recreation areas	0.00	0.00	0.00	0.00
Pastures	Forest	0.00	0.00	0.00	0.00
Agricultural areas	Arable land	14.59	0.00	0.00	0.00
Agricultural areas	Permanent crops	40.00	1.00	0.10	0.00
Agricultural areas	Pastures	0.00	0.00	0.00	0.00
Agricultural areas	Agricultural areas	275.83	0.50	0.05	0.00
Agricultural areas	Residential	47.21	27.86	2.79	0.00
Agricultural areas	Industry & commerce	58.86	3.75	0.38	0.00
Agricultural areas	Recreation areas	20.00	0.00	0.00	0.00
Agricultural areas	Forest	10.00	0.00	0.00	0.00
Residential	Arable land	0.00	0.00	0.00	0.00
Residential	Permanent crops	0.00	0.00	0.00	0.00
Residential	Pastures	0.00	0.00	0.00	0.00
Residential	Agricultural areas	0.00	1.00	0.10	0.00
Residential	Residential	1000.00	62.31	6.23	0.00
Residential	Industry & commerce	10.00	27.05	2.71	0.00
Residential	Recreation areas	0.00	1.00	0.10	0.00
Residential	Forest	0.00	0.00	0.00	0.00
Industry & commerce	Arable land	0.00	0.00	0.00	0.00
Industry & commerce	Permanent crops	0.00	0.00	0.00	0.00

Industry & commerce	Pastures	0.00	0.00	0.00	0.00
Industry & commerce	Agricultural areas	0.00	0.00	0.00	0.00
Industry & commerce	Residential	0.00	83.28	8.33	0.00
Industry & commerce	Industry & commerce	1000.00	85.41	8.54	0.00
Industry & commerce	Recreation areas	10.00	2.00	0.20	0.00
Industry & commerce	Forest	0.00	0.00	0.00	0.00
Recreation areas	Arable land	0.00	0.00	0.00	0.00
Recreation areas	Permanent crops	0.00	0.00	0.00	0.00
Recreation areas	Pastures	0.00	0.00	0.00	0.00
Recreation areas	Agricultural areas	0.00	0.00	0.00	0.00
Recreation areas	Residential	0.00	40.33	4.03	0.00
Recreation areas	Industry & commerce	0.00	38.20	3.82	0.00
Recreation areas	Recreation areas	1000.00	85.41	8.54	0.00
Recreation areas	Forest	0.00	0.00	0.00	0.00
Forest	Arable land	23.10	0.00	0.00	0.00
Forest	Permanent crops	61.80	0.00	0.00	0.00
Forest	Pastures	0.00	0.00	0.00	0.00
Forest	Agricultural areas	98.68	1.00	0.10	0.00
Forest	Residential	43.27	0.00	0.00	0.00
Forest	Industry & commerce	20.00	0.00	0.00	0.00
Forest	Recreation areas	61.30	6.89	0.69	0.00
Forest	Forest	250.00	7.20	0.72	0.00
Road & rail	Arable land	0.00	0.00	0.00	0.00
Road & rail	Permanent crops	0.00	0.00	0.00	0.00
Road & rail	Pastures	0.00	0.00	0.00	0.00
Road & rail	Agricultural areas	0.00	0.00	0.00	0.00
Road & rail	Residential	0.00	72.95	7.29	0.00
Road & rail	Industry & commerce	0.00	0.00	0.00	0.00
Road & rail	Recreation areas	0.00	2.00	0.20	0.00
Road & rail	Forest	0.00	0.00	0.00	0.00
Port area	Arable land	0.00	0.00	0.00	0.00
Port area	Permanent crops	0.00	0.00	0.00	0.00
Port area	Pastures	0.00	0.00	0.00	0.00
Port area	Agricultural areas	0.00	0.00	0.00	0.00
Port area	Residential	0.00	0.00	0.00	0.00
Port area	Industry & commerce	0.00	0.00	0.00	0.00
Port area	Recreation areas	10.00	76.39	7.64	0.00
Port area	Forest	0.00	0.00	0.00	0.00
Airports	Arable land	0.00	0.00	0.00	0.00
Airports	Permanent crops	0.00	0.00	0.00	0.00
Airports	Pastures	0.00	0.00	0.00	0.00
Airports	Agricultural areas	0.00	0.00	0.00	0.00
Airports	Residential	0.00	0.00	0.00	0.00
Airports	Industry & commerce	0.00	0.00	0.00	0.00
Airports	Recreation areas	0.00	0.00	0.00	0.00
Airports	Forest	0.00	0.00	0.00	0.00
Mine & dump sites	Arable land	0.00	0.00	0.00	0.00
Mine & dump sites	Permanent crops	0.00	0.00	0.00	0.00
Mine & dump sites	Pastures	0.00	0.00	0.00	0.00
Mine & dump sites	Agricultural areas	0.00	18.85	1.88	0.00
Mine & dump sites	Residential	0.00	1.00	0.10	0.00
Mine & dump sites	Industry & commerce	0.00	1.00	0.10	0.00
Mine & dump sites	Recreation areas	0.00	2.00	0.20	0.00
Mine & dump sites	Forest	0.00	0.00	0.00	0.00
Fresh water	Arable land	0.00	0.00	0.00	0.00



Fresh water	Permanent crops	0.00	0.00	0.00	0.00
Fresh water	Pastures	0.00	80.15	8.01	0.00
Fresh water	Agricultural areas	0.00	0.00	0.00	0.00
Fresh water	Residential	0.00	0.00	0.00	0.00
Fresh water	Industry & commerce	0.00	0.00	0.00	0.00
Fresh water	Recreation areas	0.00	0.00	0.00	0.00
Fresh water	Forest	0.00	0.00	0.00	0.00
Marine water	Arable land	0.00	0.00	0.00	0.00
Marine water	Permanent crops	0.00	0.00	0.00	0.00
Marine water	Pastures	0.00	0.00	0.00	0.00
Marine water	Agricultural areas	0.00	0.00	0.00	0.00
Marine water	Residential	0.00	0.00	0.00	0.00
Marine water	Industry & commerce	0.00	0.00	0.00	0.00
Marine water	Recreation areas	0.00	0.00	0.00	0.00
Marine water	Forest	0.00	0.00	0.00	0.00

Table 3.E4. Final neighbourhood rule parameter values, Madrid

From	To	Distance			
		0	1	2	5
		Influence			
Natural areas	Arable land	15.40	2.75	0.28	0.00
Natural areas	Permanent crops	60.49	0.00	0.00	0.00
Natural areas	Pastures	0.00	0.00	0.00	0.00
Natural areas	Agricultural areas	20.00	0.00	0.00	0.00
Natural areas	Residential	20.00	0.00	0.00	0.00
Natural areas	Industry & commerce	32.12	0.00	0.00	0.00
Natural areas	Recreation areas	19.85	2.75	0.28	0.00
Natural areas	Forest	33.13	2.13	0.21	0.00
Arable land	Arable land	250.00	7.70	0.77	0.00
Arable land	Permanent crops	40.00	2.75	0.28	0.00
Arable land	Pastures	0.00	0.00	0.00	0.00
Arable land	Agricultural areas	28.37	2.75	0.28	0.00
Arable land	Residential	52.48	2.75	0.28	0.00
Arable land	Industry & commerce	69.50	2.13	0.21	0.00
Arable land	Recreation areas	38.70	0.00	0.00	0.00
Arable land	Forest	81.46	0.00	0.00	0.00
Permanent crops	Arable land	10.00	0.00	0.00	0.00
Permanent crops	Permanent crops	572.17	21.48	2.15	0.00
Permanent crops	Pastures	0.00	0.00	0.00	0.00
Permanent crops	Agricultural areas	0.00	0.00	0.00	0.00
Permanent crops	Residential	0.00	0.00	0.00	0.00
Permanent crops	Industry & commerce	0.00	0.00	0.00	0.00
Permanent crops	Recreation areas	0.00	0.00	0.00	0.00
Permanent crops	Forest	0.00	0.00	0.00	0.00
Pastures	Arable land	0.00	0.00	0.00	0.00
Pastures	Permanent crops	0.00	0.00	0.00	0.00
Pastures	Pastures	735.59	10.00	1.00	0.00
Pastures	Agricultural areas	0.00	0.00	0.00	0.00
Pastures	Residential	0.00	2.75	0.28	0.00
Pastures	Industry & commerce	0.00	0.00	0.00	0.00
Pastures	Recreation areas	0.00	0.00	0.00	0.00
Pastures	Forest	0.00	0.00	0.00	0.00
Agricultural areas	Arable land	10.00	0.00	0.00	0.00
Agricultural areas	Permanent crops	20.00	8.51	0.85	0.00
Agricultural areas	Pastures	0.00	0.00	0.00	0.00
Agricultural areas	Agricultural areas	353.33	9.33	0.93	0.00
Agricultural areas	Residential	42.96	0.00	0.00	0.00
Agricultural areas	Industry & commerce	52.28	0.00	0.00	0.00
Agricultural areas	Recreation areas	63.12	0.00	0.00	0.00
Agricultural areas	Forest	0.00	0.00	0.00	0.00
Residential	Arable land	0.00	0.00	0.00	0.00
Residential	Permanent crops	0.00	0.00	0.00	0.00
Residential	Pastures	0.00	0.00	0.00	0.00
Residential	Agricultural areas	0.00	0.00	0.00	0.00
Residential	Residential	1000.00	61.30	6.13	0.00
Residential	Industry & commerce	10.00	35.57	3.56	0.00
Residential	Recreation areas	10.00	2.75	0.28	0.00
Residential	Forest	0.00	0.00	0.00	0.00
Industry & commerce	Arable land	0.00	2.75	0.28	0.00
Industry & commerce	Permanent crops	0.00	0.00	0.00	0.00

Industry & commerce	Pastures	0.00	0.00	0.00	0.00
Industry & commerce	Agricultural areas	0.00	0.00	0.00	0.00
Industry & commerce	Residential	0.00	3.44	0.34	0.00
Industry & commerce	Industry & commerce	1000.00	55.73	5.57	0.00
Industry & commerce	Recreation areas	0.00	10.33	1.03	0.00
Industry & commerce	Forest	0.00	0.00	0.00	0.00
Recreation areas	Arable land	0.00	0.00	0.00	0.00
Recreation areas	Permanent crops	0.00	0.00	0.00	0.00
Recreation areas	Pastures	0.00	0.00	0.00	0.00
Recreation areas	Agricultural areas	0.00	0.00	0.00	0.00
Recreation areas	Residential	0.00	5.57	0.56	0.00
Recreation areas	Industry & commerce	0.00	0.00	0.00	0.00
Recreation areas	Recreation areas	500.00	20.16	2.02	0.00
Recreation areas	Forest	0.00	0.00	0.00	0.00
Forest	Arable land	10.00	0.00	0.00	0.00
Forest	Permanent crops	0.00	0.00	0.00	0.00
Forest	Pastures	0.00	0.00	0.00	0.00
Forest	Agricultural areas	0.00	0.00	0.00	0.00
Forest	Residential	0.00	0.00	0.00	0.00
Forest	Industry & commerce	0.00	0.00	0.00	0.00
Forest	Recreation areas	10.00	0.00	0.00	0.00
Forest	Forest	1000.00	5.57	0.56	0.00
Road & rail	Arable land	0.00	0.00	0.00	0.00
Road & rail	Permanent crops	0.00	0.00	0.00	0.00
Road & rail	Pastures	0.00	0.00	0.00	0.00
Road & rail	Agricultural areas	0.00	0.00	0.00	0.00
Road & rail	Residential	0.00	0.00	0.00	0.00
Road & rail	Industry & commerce	0.00	68.19	6.82	0.00
Road & rail	Recreation areas	0.00	0.00	0.00	0.00
Road & rail	Forest	0.00	0.00	0.00	0.00
Port area	Arable land	0.00	0.00	0.00	0.00
Port area	Permanent crops	0.00	0.00	0.00	0.00
Port area	Pastures	0.00	0.00	0.00	0.00
Port area	Agricultural areas	0.00	0.00	0.00	0.00
Port area	Residential	0.00	0.00	0.00	0.00
Port area	Industry & commerce	0.00	0.00	0.00	0.00
Port area	Recreation areas	0.00	0.00	0.00	0.00
Port area	Forest	0.00	0.00	0.00	0.00
Airports	Arable land	0.00	0.00	0.00	0.00
Airports	Permanent crops	0.00	0.00	0.00	0.00
Airports	Pastures	0.00	0.00	0.00	0.00
Airports	Agricultural areas	0.00	0.00	0.00	0.00
Airports	Residential	0.00	0.00	0.00	0.00
Airports	Industry & commerce	0.00	19.35	1.93	0.00
Airports	Recreation areas	0.00	0.00	0.00	0.00
Airports	Forest	0.00	0.00	0.00	0.00
Mine & dump sites	Arable land	0.00	0.00	0.00	0.00
Mine & dump sites	Permanent crops	0.00	31.31	3.13	0.00
Mine & dump sites	Pastures	0.00	0.00	0.00	0.00
Mine & dump sites	Agricultural areas	0.00	0.00	0.00	0.00
Mine & dump sites	Residential	0.00	51.47	5.15	0.00
Mine & dump sites	Industry & commerce	0.00	11.15	1.11	0.00
Mine & dump sites	Recreation areas	10.00	11.15	1.11	0.00
Mine & dump sites	Forest	0.00	0.00	0.00	0.00
Fresh water	Arable land	0.00	0.00	0.00	0.00

Fresh water	Permanent crops	0.00	0.00	0.00	0.00
Fresh water	Pastures	0.00	0.00	0.00	0.00
Fresh water	Agricultural areas	0.00	0.00	0.00	0.00
Fresh water	Residential	0.00	0.00	0.00	0.00
Fresh water	Industry & commerce	0.00	0.00	0.00	0.00
Fresh water	Recreation areas	0.00	0.00	0.00	0.00
Fresh water	Forest	0.00	0.00	0.00	0.00
Marine water	Arable land	0.00	0.00	0.00	0.00
Marine water	Permanent crops	0.00	0.00	0.00	0.00
Marine water	Pastures	0.00	0.00	0.00	0.00
Marine water	Agricultural areas	0.00	0.00	0.00	0.00
Marine water	Residential	0.00	0.00	0.00	0.00
Marine water	Industry & commerce	0.00	0.00	0.00	0.00
Marine water	Recreation areas	0.00	0.00	0.00	0.00
Marine water	Forest	0.00	0.00	0.00	0.00

## **4 A hybrid (semi) automatic calibration method for Cellular Automata land-use models: Combining evolutionary algorithms with process understanding**

Newland, C.P., Maier, H.R., Zecchin, A.C., Newman, J.P. & Van Delden, H. A hybrid (semi) automatic calibration method for Cellular Automata land-use models: Combining evolutionary algorithms with process understanding. (Adapted for consistency).

## Statement of Authorship

Title of Paper	A hybrid (semi) automatic calibration method for Cellular Automata land-use models: Combining evolutionary algorithms with process understanding		
Publication Status	<input type="checkbox"/> Published	<input type="checkbox"/> Accepted for publication	
	<input type="checkbox"/> Submitted for publication	<input checked="" type="checkbox"/> Unpublished and Un-submitted work written in manuscript style	
Publication Details			

## Principal Author

Name of Principal Author (Candidate)	Charles P. Newland		
Contribution to the Paper	Designed scope of study and experimental procedure, developed software, conducted experiments, performed analysis of results. Prepared manuscript.		
Overall percentage (%)	70%		
Certification	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.		
Signature		Date	28 <sup>th</sup> Feb 2018

## Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

- i. The candidate's stated contribution to the publication is accurate (as detailed above);
- ii. Permission is granted for the candidate to include the publication in the thesis; and
- iii. The sum of all co-author's contributions is equal to 100% less the candidate's stated contribution.

Name of Co-author	Holger R. Maier		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript.		
Signature		Date	28 <sup>th</sup> Feb. 2018

Name of Co-author	Aaron C. Zecchin		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript.		
Signature		Date	28 <sup>th</sup> Feb. 2018

Name of Co-author	Jeffrey P. Newman		
Contribution to the Paper	Assisted with software development for super-computer and conducting experiments. Reviewed manuscript.		
Signature		Date	28 <sup>th</sup> Feb. 2018

Name of Co-author	Hedwig Van Delden		
Contribution to the Paper	Assisted with developing scope of study, experimental procedure, and analysis of results. Reviewed manuscript.		
Signature		Date	28 <sup>th</sup> Feb. 2018

## **A hybrid (semi) automatic calibration method for Cellular Automata land-use models: Combining evolutionary algorithms with process understanding**

### **Abstract**

This paper presents a hybrid automatic calibration method for transition potential based Cellular Automata land-use models that integrates process-specific and optimisation-based calibration methods. Process-specific methods are computationally efficient and result in calibration parameters that align with process understanding, while optimisation-based approaches are able to identify parameters that provide the optimal trade-offs between calibration objectives, but are computationally demanding and less likely to generate parameters that are consistent with process understanding. To address the shortcomings of these two types of existing approaches, the proposed hybrid approach uses the outcomes of process-specific methods to initiate formal optimisation in promising regions of the parameter space, increasing computational efficiency and alignment with process understanding. The utility of the approach is tested via a case study application to Madrid, Spain, and was found to outperform both existing approaches in terms of objective performance and solution quality, generating simulated output maps and parameters more consistent with process understanding.



## 4.1 Introduction

Land-use models are being used increasingly in a policy setting to support environmental and urban planning (Van Delden et al., 2011). Such models are used to explore the impacts of policy alternatives across a diverse range of disciplines, including river basin management (Van Delden et al., 2007), regional growth planning (Rutledge et al., 2008), development of sustainable agricultural practises (Wickramasuriya et al., 2009), and urban forecasting (Chaudhuri and Clarke, 2013a, Berberoğlu et al., 2016). To effectively model land-use changes, the complex ecological and socio-economic drivers must be effectively captured (Lambin et al., 2001). A preferred approach for modelling land-use changes is to use Cellular Automata (CA) based land-use models, due to their simplicity (Santé et al., 2010) and ability to explore land-use dynamics of large areas without detailed data requirements (Hewitt et al., 2014).

The intuitiveness and effectiveness of Land Use Cellular Automata (LUCA) models has led to the development of the generic modelling frameworks SLEUTH (Clarke et al., 1997) and Metronamica (Van Delden and Hurkens, 2011) that have well-tested model architectures, which has facilitated the application of such models to many different case studies (Chaudhuri and Clarke, 2013b, Van Delden et al., 2011). With generic platforms essentially obviating model software development requirements (Hewitt et al., 2014), there has been a greater research focus on the calibration of LUCA models (Blecic et al., 2015, Clarke-Lauer and Clarke, 2011, Li et al., 2013, Van Vliet et al., 2013b, Şalap-Ayça et al., 2018, Newland et al., 2018a, Newland et al., 2018b).

The process of calibrating LUCA models, which involves initial parameter setting, iterative parameter adjustment, and selection of a final parameter set (Newland et al., 2018a), is a complex process, because land-use change is a path dependent process with an uncertain outcome that must capture multiple interdependent drivers to accurately replicate land-use change dynamics (Brown et al., 2005). Conventionally, calibration has been a manual procedure, utilising the modeller's process understanding to address these issues (White et al., 1997, Barredo et al., 2003). However, manual calibration of LUCA models is time consuming (García et al., 2013) and inherently subjective (Jafarnezhad et al., 2016), making it a difficult process to repeat. Hence, to make the process more repeatable and efficient, there has been a large focus on automating LUCA model calibration.

The automation of LUCA model calibration requires quantitative measures of model performance that enable the simulated model output (i.e. a land-use map) to be compared with the corresponding data. Due to the inherent path-dependence of land-use change as a process, two separate properties of model performance must be considered, including locational agreement and landscape pattern structure (Newland et al., 2018a). Performance metrics quantifying these two classes of model performance assessment have been incorporated into automatic calibration methods for LUCA models that are implemented with a single dynamic urban class (Silva and Clarke, 2002, Wu, 2002, Li et al., 2013, Clarke, 2018) and those that consider multiple dynamic land-use classes (Blecic et al., 2015, García et al., 2013, Straatman et al., 2004, Van Vliet et al., 2013b), which are most commonly transition potential models derived from White and Engelen (1993c). The consideration of multiple dynamic land-use classes invariably makes the calibration of LUCA models more complex, as such models possess a significantly larger number of parameters (often on the order of hundreds) that require calibration to values that should be consistent with process understanding.

Previous attempts at (semi) automatic calibration of LUCA models that consider multiple dynamic land-use classes can be broadly divided into two types. The first approach, termed process-specific, focuses on efficient calibration by exploiting process knowledge and mathematical properties of LUCA models to generate a set of parameters underlying a certain model process that is consistent with process understanding within a reasonable computational budget, achievable using a desktop PC (Straatman et al., 2004, Maas et al., 2005, Van Vliet et al., 2013b, Newland et al., 2018b). This efficiency focus means there is limited exploration of the impact of the trade-off between the objectives mentioned previously, limiting the ability to identify the best possible set of model parameters.

The second approach to automate calibration is to use an optimisation-based approach, which aims to comprehensively search the LUCA model parameter space to identify the set (or sets in the case of multi-objective optimisation) of parameters that optimise the measures of model performance, and is a common approach to automatically calibrate LUCA models (Clarke, 2018, Li et al., 2013, Liao et al., 2014). Both single-objective (García et al., 2013, Blecic et al., 2015) and multi-objective (Newland et al., 2018a) optimisation algorithms have been used, with multi-objective optimisation allowing for a detailed exploration of the trade-off between the model performance objectives, and the influence this has on the resultant model parameters. However, the black-box nature of optimisation means the resultant parameters obtained may not be consistent with process understanding. Also, parallel computing resources are often

required for practical implementation of optimisation-based approaches (Blecic et al., 2015, Newland et al., 2018a).

Both of the approaches to (semi) automatic calibration of transition potential LUCA models mentioned above have advantages and disadvantages. Process-specific methods are computationally efficient, and result in calibration parameters that align with process understanding, but have limited ability to explore different model parameterisations, making it unlikely the best parameter set is identified. Optimisation-based approaches have the ability to explore a large number of different possible model parameterisations, but are computationally demanding, often requiring parallel computing resources for practical implementation, and are less likely to generate parameters consistent with process understanding. To address the shortcomings of both methods, this paper proposes a hybrid approach that builds on the strengths of both methods. Therefore the objectives of this paper are: (i) to introduce a method for integrating the results of a process-specific automatic calibration approach into an optimisation-based approach; and (ii) to evaluate the utility of the proposed hybrid approach, and compare the results with standard implementations of the other two approaches. The remainder of this paper is organised as follows: Section 4.2 describes the proposed hybrid automatic calibration approach. Section 4.3 describes the case study application of the proposed approach, and Section 4.4 presents and discusses the results. The conclusions and recommendations following this body of work are presented in Section 4.5.

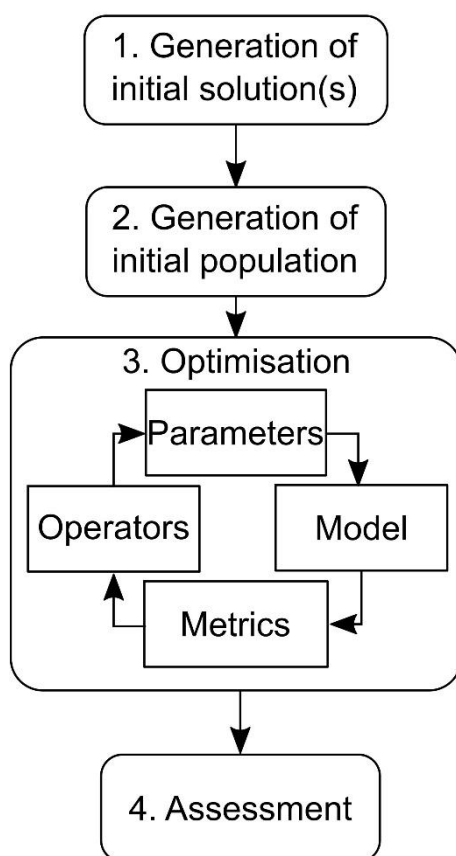
#### **4.2 Proposed approach to integrating process understanding into optimisation-based automatic calibration**

The proposed approach, shown in Figure 4.1, aims to improve the efficiency and parameter validity of multi-objective metaheuristic optimisation (e.g. genetic algorithm) automatic calibration approaches by seeding the optimisation process in promising regions of the parameter space, which is achieved with the assistance of a process-specific (or possibly manual) calibration method. This is likely to improve the speed with which globally optimal solutions are identified, therefore addressing one of the major shortcomings of optimisation-based calibration methods. Additionally, by starting the optimisation in regions of the parameter space that align with domain knowledge, there is an increased likelihood that the parameters identified at the end of the optimisation process are more physically plausible, addressing the other major potential drawback of optimisation-based calibration. The approach is tailored to metaheuristics because of their demonstrated ability to automatically calibrate LUCA models (Veerbeek et al., 2015, Blecic et al., 2015, García et al., 2013, Clarke, 2018).

When metaheuristic optimisation approaches are used, the optimisation process (shown in Figure 4.1) involves the generation of populations of solutions consisting of sets of decision variable combinations (in this case model parameters), which are input into a model to generate simulated output, with the performance of each member of the population quantified by the objective function (in this case LUCA model performance metrics). Based on the optimisation operators (which vary depending on the algorithm used) new sets of decision variables are generated, and the process is repeated until some stopping criteria are met.

Traditionally, a metaheuristic generates an initial population of candidate solutions using a random sampling method. However, the random generation of initial candidate solutions can impact the convergence rate of the metaheuristic, and potentially the quality of the final outcome obtained. This has been addressed in other fields by using domain knowledge to generate an initial population of solutions to “hot-start” the optimisation process, and has the ability to improve both the efficiency of the optimisation process, and the plausibility of the obtained results (Kang and Lansey, 2012, Bi et al., 2015, Bi et al., 2016a). Consequently, this philosophy has been adopted here by introducing an approach that generates the initial population of solutions for the metaheuristic optimisation algorithm with the aid of solutions identified using either a process-specific (semi) automatic (or possibly manual) calibration method.

Given the multi-objective nature of the metaheuristic optimisation process, a number of Pareto optimal parameter sets (solutions) are generated, the relative merits of which can only be determined with the aid of user input. Consequently, the final step of the proposed approach involves model assessment, where parameter sets on the Pareto front are scrutinised in terms of the consistency of the simulated land-use maps and calibrated parameter sets with process understanding. In addition, the calibrated model is assessed with the aid of performance benchmarks and independent validation data. Details of each of these steps are given below.



*Figure 4.1. Conceptual framework for proposed approach to integrating process understanding into optimisation-based automatic calibration*

#### **4.2.1 Generation of initial solution**

For transition potential LUCA models, the generation of initial, high-quality solutions can be performed using a process-specific (semi) automatic calibration method (e.g. Straatman et al., 2004, Maas et al., 2005, Van Vliet et al., 2013b), which will generate a set of calibrated model parameters that are consistent with process knowledge in an efficient manner. The method considered most appropriate is case study and user dependent. Alternatively, such initial solutions can also be obtained using manual calibration.

#### **4.2.2 Initial population generation**

The purpose of this step is to create an initial population of solutions that is located in promising regions of the parameter space, whilst also having sufficient diversity to enable exploration of other regions, which increases the chance that the globally optimal solutions are identified. Hence, the initial population not only includes the solutions that occur in good regions of the parameter space identified while generating the initial solution (Section 4.2.1), but also solutions in the vicinity of these solutions, as well as randomly generated solutions, as has been

done in other problem domains (Bi et al., 2015, Bi et al., 2016b). Consequently, the initial population is composed of three distinct types of solutions:

1. Randomly generated solutions, following the standard initialisation procedure for a population-based metaheuristic;
2. Process-specific solutions, generated using a process-specific (semi) automatic (or possibly manual) calibration approach; and
3. Sampled solutions, generated by sampling from a distribution centred on the parameters of the process-specific solution.

The first set of solutions is included to ensure sufficient diversity in the initial population, whereas the second and third types of solutions are used to initialise the optimisation in promising regions of the parameter space. What fraction of the total population is made up of the different solution types is user defined, and will vary depending on the case-study, and the desired balance between convergence and exploration of the parameter space. Hence, the user must determine what percentage of the initial population is comprised of randomly generated solutions, process-specific solutions, and sampled solutions, and the required number of samples from each category has to be generated accordingly.

As mentioned above, the sampled solutions are used to commence the optimisation process in a promising region of the parameter space (i.e. in the vicinity of the initially generated solution), by producing solutions that are similar to the one obtained from the process-specific or manual calibration approach. As suggested by Bi et al. (2016a), this is achieved by fitting a triangular distribution to each parameter, centred on the process-specific value obtained, as a triangular distribution permits sufficient differentiation of the parameter value within a fixed range (as opposed to a normal distribution, where the range is unbounded).

An example of fitting a triangular distribution to a model parameter is shown in Figure 4.2, where the mean of the distribution  $\theta$ , shown by the vertical grey line, corresponds to a process-specific parameter value obtained in the initial calibration stage (Section 4.2.1), and the upper and lower limits, expressed as a percentage of the obtained parameter value, influencing the range of the potential sampled values. As shown, by increasing the limits, the potential range of parameter values obtained is broader. It is important to use limits that ensure sufficient variation of the parameters without generating values that are inconsistent with process understanding.




Figure 4.2. Comparison of triangular distributions with varying limits for sampling of parameter values

### 4.2.3 Optimisation

The proposed optimisation approach is an adapted version of the multi-objective optimisation framework of Newland et al. (2018a), where, rather than starting the optimisation process with a random population of solutions, the initial population is made up of a mixture of random and “good” initial solutions, obtained using the process outlined in Section 4.2.2. As part of the optimisation process, the optimisation algorithm takes the initial population of LUCA model parameter combinations and iteratively improves them by running each through the LUCA model to generate a simulated output map and calculating a set of land-use metrics to quantify the different aspects of LUCA model performance (i.e. locational agreement and landscape pattern structure) for each solution (Figure 4.1). The population of sets of parameter values is then adjusted based on feedback of the relative performance of the different parameter sets in the population using the metaheuristic operators (e.g. selection, cross-over and mutation in the case of a genetic algorithm). This process is repeated until certain stopping criteria, such as a fixed number of generations or convergence as indicated by performance measures such as the hyper-volume, are met. Given that locational agreement and landscape pattern structure are

competing optimisation objectives (Trunfio, 2006, Newland et al., 2018a), the outcome of the optimisation process is not a single set of optimal LUCA model parameters, but a set of parameters that represent an optimal trade-off between the two objectives, called the Pareto front.

#### 4.2.4 Assessment

The last step is to assess the final output obtained. This assessment should consider the objective performance of the final output, the metrics for the calibration period evaluated against benchmark models of land-use change (Hagen-Zanker and Lajoie, 2008) and objective performance for an independent validation period. Assessment should also consider the physical plausibility of the output, by evaluating whether the simulated output maps and model parameters are consistent with process knowledge. After the assessment has been conducted, a final calibrated model can be determined for long-term scenario analysis.

### 4.3 Case study

To evaluate the utility of the proposed hybrid method, the case study implementation detailed in the following section is used, with details given for each step in the proposed approach. The Metronamica land-use model is used, with a case study of Madrid, Spain. Metronamica is a generic LUCA model that uses a transition potential as the allocation mechanism for land-use changes (Van Delden and Hurkens, 2011), and has had numerous applications to diverse regions (Aljoufie et al., 2016, Furtado et al., 2012, Van Delden et al., 2010, Shi et al., 2012). The transition potential is a function of the modelled processes considered, given in Equation 4.1:

$$TP_{c,k} = \gamma \cdot A_{c,k} \cdot S_{c,k} \cdot N_{c,k} \cdot Z_{c,k} \quad (4.1)$$

where  $TP_{c,k}$  is the transition potential for land-use class  $k$  in cell  $c$  (note that all subscripts have the same interpretation),  $A_{c,k}$  is the accessibility (the influence of provisions of infrastructure),  $S_{c,k}$  is the suitability (the influence of bio-physical factors such as slope),  $N_{c,k}$  is the neighbourhood effect (the spatial interactions between different land-uses in the competition for space),  $Z_{c,k}$  is the zoning (the influence of spatial planning) and  $\gamma$  is a stochastic element, included to capture the uncertainty of human decisions. Additional information about the Metronamica model and the parameterisation of the processes included is available from other sources (RIKS, 2015, Engelen and White, 2008, van Delden and Vanhout, 2018).

The Madrid case study is developed using the CORINE land-use data set (Haines-Young et al., 2006). The model uses data from 1990 to 2000 as the calibration period, and 2000 to 2006 as



the validation period. A land-use map of the Madrid region in 1990 is shown in Figure 4.3. It covers an area of 400 by 400 cells at a resolution of 250 metres, covering a total area of 10,000 km<sup>2</sup>. The 48 CORINE level 3 land-use classes are reclassified to the 14 classes shown in Figure 4.3. Of the 14 classes, 1 is passive, 8 are actively modelled, and 5 are static classes (that do not change throughout the simulation). The case study includes major roads data for accessibility.

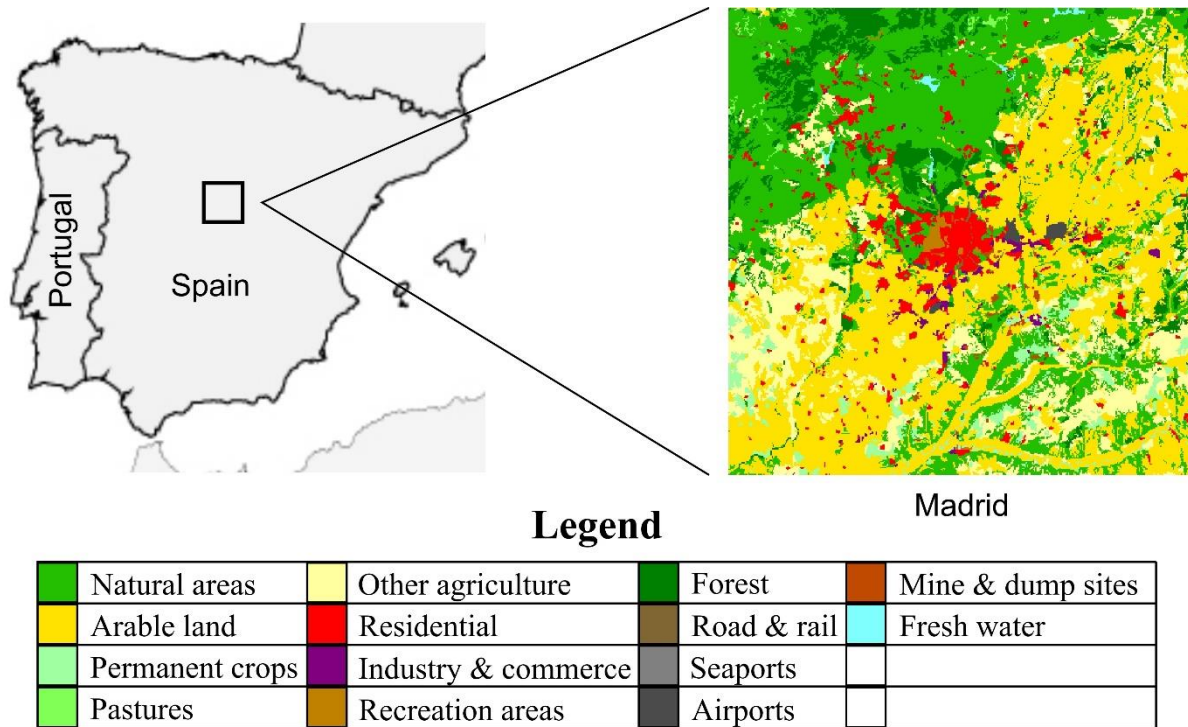


Figure 4.3. Location of Madrid region and rasterised land-use map for 1990

As the case study is based in an established urban region, the main drivers of land-use changes are likely to relate to the growth of existing socio-economic land-uses, with increased expansion of urban cores driven by self-organising behaviour (Couclelis, 1989, Batty and Longley, 1994, White and Engelen, 1993a). Hence, the most important processes to consider are those driving the expansion of urban regions, the neighbourhood interactions and accessibility (Verburg et al., 2004). Consequently, suitability and zoning are not included for calibration in this case study. The parameters for neighbourhood rules and accessibility are selected for automatic calibration, and each neighbourhood rule is parameterised as shown in Figure 4.4, following the form:

$$y(x) = \begin{cases} c & \text{for } x = 0 \\ ae^{-bx} & \text{for } 0 < x \leq x_c \\ 0 & \text{for } x > x_c \end{cases} \quad (4.2)$$

where  $a$  and  $b$  are the controlling parameters of the neighbourhood rule,  $x$  is the distance,  $y(x)$  is the influence value,  $c$  is the locus point of inertia or conversion, and  $x_c$  is the critical distance where the influence is set to zero. As there are 8 active classes and 14 total classes, there are a total of 112 neighbourhood rules, and hence 336 neighbourhood rule parameters for calibration. For accessibility, two parameters are included for each actively modelled land-use class to capture the influence of major roads, meaning a total of 16 accessibility parameters are calibrated. Hence, in total, 352 parameters require calibration.




Figure 4.4. Parameterisation of neighbourhood rule using an exponential decay function

For the (semi) automatic calibration methods, objectives of locational agreement and landscape pattern structure are required. Locational agreement is quantified using two metrics: Fuzzy Kappa (Hagen-Zanker, 2009) and Fuzzy Kappa Simulation (Van Vliet et al., 2013b). Two metrics are used to balance between the two types of agreement these metrics quantify, that of all cells in the land-use map, measured by Fuzzy Kappa (FK), and cells that transition between the two time slices, measured by Fuzzy Kappa Simulation (FKS). Both metrics have the same functional form:

$$FK = \frac{P_O - P_E}{1 - P_E} \quad (4.3)$$

where  $P_O$  is the observed agreement, and  $P_E$  is the expected agreement between the two categorical raster maps. FK and FKS vary by how the observed and expected agreement is calculated.

To measure landscape pattern structure, the clumpiness metric is used (McGarigal, 2014). As clumpiness is measured at the class level and a single value is required as an objective, the error between the actively allocated land-use classes is aggregated to a single value by taking the area-weighted average for the eight actively modelled land-use classes, henceforth referred to as the Area-Weighted Clumpiness Error (AWCE):

$$AWCE = \frac{1}{N_A} \sum_i^A CE_i \cdot n_i \quad (4.4)$$

where  $N_A$  is the total number of cells that contain actively allocated land-use classes,  $A$  is the total number of actively modelled land-use classes,  $CE_i$  is the clumpiness error of active class  $i$ , and  $n_i$  is the number of cells occupied by the class  $i$ .

### 4.3.1 Implementation of proposed approach

#### 4.3.1.1 Generation of initial solution

To generate an initial candidate solution that is consistent with process understanding in an efficient manner, the semi-automatic calibration method of Newland et al. (2018b) is used. This method uses empirical analysis to evaluate the different types of spatial interactions, which can be organised into four interaction groups: inertia points (*IP*), conversion points (*CP*), self-influence tails (*ST*) and cross-influence tails (*CT*), to identify the major set of interactions driving land-use changes in a region, and efficiently calibrates the corresponding neighbourhood rules. This method consists of four stages:

1. Interaction elimination;
2. Parameter categorisation;
3. Coarse parameter adjustment; and
4. Fine parameter adjustment.

A brief summary of each stage is provided below.

Stage 1, the **interaction elimination stage**, is used to reduce the complexity of the calibration problem by identifying the main set of interactions within each interaction group that is driving land-use changes within the region of interest. The parameters corresponding to these rules are subsequently calibrated and the parameters that do not correspond to major interactions are eliminated from consideration (i.e. set to zero). To achieve this, empirical analysis and significance testing are applied to the calibration data to identify the meaningful interactions within the region.

For the specific case study, all inertia points and self-influence tails are included for calibration, and only conversion points and cross-influence tails are eliminated. Significance testing is performed using the Mann-Whitney U-test to generate a z-score for the conversion points and cross-influence tails, by evaluating the presence of a certain land-use class in the neighbourhood of cells that transitioned to another land-use class, relative to the presence of the neighbourhood class in the entire land-use map. Empirical analysis of the conversion points is based on the contingency table, a form of confusion matrix (Congalton, 1991) that logs the land-use class of each cell for two time slices, from which the conversion rate is calculated. Conversion points corresponding to a conversion rate of  $\geq 2.5\%$  and a significance level of  $> 1.96$  (the 95% confidence limit) are included for calibration. The enrichment factor is used to perform the empirical analysis of the cross-influence tails, to identify meaningful attractive influences. Cross-influence tails corresponding to log-scaled enrichment factor values of greater than zero, suggesting over-representation and hence an attractive influence, and a significance level  $> 1.96$  are included for calibration.

Stage 2, the **parameter categorisation stage**, is used to further reduce the complexity of the calibration problem. This is achieved first by introducing a set of meta-parameters,  $\theta_{CP}$ ,  $\theta_{ST}$ , and  $\theta_{CT}$ , which express the *inter-type* importance of each interaction relative to the inertia point interactions (e.g. how important conversion points are compared to inertia points). Second, parameter categorisation limits the parameters' values to a finite set of possible categorised values, and allocates each parameter included for calibration within each interaction group to one category value based on empirical analysis. Hence, the *intra-type* importance for each interaction type is defined. The categorisation is performed by using empirical analysis to assign each parameter within the four parameter groups to one of a user-defined number of categories. For this case study, three categories are used for each parameter group, graded as either low, medium or high. Allocation of the inertia and conversion points is based on the inertia and conversion rate, respectively, both derived from the contingency table. Allocation of the self-influence and cross-influence tails is based on the enrichment factor.

Stage 3, the **coarse parameter adjustment stage**, is where LUCA model performance is improved via calibration of the meta-parameters (introduced in Stage 2), which control the categorised neighbourhood weighting parameters. To perform this, one meta-parameter is sampled within a range, and the meta-parameter value that optimises model performance is selected, based on an evaluation of the performance metrics (FK, FKS and AWCE) and discursive interpretation of the output. This meta-parameter is then fixed, and the next meta-

parameter sampled, until all three have been tuned. The tuned meta-parameters are used to generate the set of neighbourhood rules that serve as the input to the fine parameter adjustment stage. For this implementation, the meta-parameters were sampled in the following order:  $\theta_{ST}$ ,  $\theta_{CP}$ , and  $\theta_{CT}$ .

Stage 4, the **fine parameter adjustment stage**, aims to further improve LUCA model performance. The fine parameter adjustment is initialised using the output of the coarse adjustment stage, corresponding to a good starting point in the parameter space. The fine parameter adjustment considers the individual neighbourhood weighting parameters, iteratively refining each parameter to optimise the performance metrics within the available computational budget, which is achieved using a line-search algorithm. As a line-search algorithm is used, the measures of calibration performance must be combined into a single measure of performance. This is achieved by taking a weighted sum of the performance objectives, which requires a range and weight to be specified for each metric considered, allowing for a preference for a certain objective. To improve the diversity of the initial population generated, two candidate solutions are used. Both use a balanced preference of objectives, however, one results in slightly better performance for FK and FKS, and the other for AWCE. The final output of this stage are two parameter sets that result in good objective performance (i.e. outperforms the performance benchmarks) and have parameters that are consistent with process understanding.

#### *4.3.1.2 Initial population generation*

The initial population generation requires the population size to be specified. The population size used is 383, as for optimisation problems with many decision variables (which is the case for this problem), a population size is required that is larger than the number of decision variables (parameters) to find non-dominated solutions (Wang et al., 2015). The population size is also informed by the parallelisation scheme adopted for optimisation (see below). The proportions of the three types of solutions (i.e. randomly generated, process-specific, sampled, see Section 4.2.2) within this population are summarised in Table 4.1, and were determined from preliminary testing. As shown, a majority of the solutions are randomly generated to ensure sufficient diversity within the search. This is because elitism used in the multi-objective evolutionary algorithm used (NSGA-II, see Section 4.3.1.3) preserves the best performing solutions, and not including a large percentage of randomly generated solutions types causes premature convergence, making exploration entirely dependent upon the probability of mutation. Two copies of the two process-specific solutions are included to explore different

promising regions of the parameter space, and the remaining percentage of the total population is generated via the sampling approach, applied to both process-specific solutions, to ensure sufficient exploration of the promising regions of the parameter space that have both good objective performance and are consistent with process knowledge.

*Table 4.1. Proportion of solution types used for initial population*

<b>Solution type</b>	<b>Number of total population</b>	<b>Percentage of total population (%)</b>
Randomly generated	304	79.4%
Process-specific	4 (2 copies of each)	1.0%
Sampled	75	19.6%

To generate the sampled solutions a triangular distribution is fitted to each neighbourhood rule parameter centred on the parameter value obtained at the conclusion of the process-specific method. Parameters that were eliminated as part of the process-specific method are set to zero, and not sampled, though still included in the subsequent optimisation to allow for further exploration of the parameter space. Different limits are trialled to ensure sufficient variation of the parameters without compromising the quality of the solution obtained, with a width of 25% of the obtained parameter value working best to achieve this. For each candidate solution, 150 sampled solutions are generated. From the combined 300 solutions, a total of 75 are extracted, based on Pareto dominance (i.e. extracting non-dominated solutions from the available candidates), to serve as the sampled input.

#### *4.3.1.3 Optimisation*

As mentioned previously, this work uses the Non-dominated Sorting Genetic Algorithm II (NSGA-II) proposed by Deb et al. (2002), because of its demonstrated ability to tune LUCA model parameters (Trunfio, 2006, Cao et al., 2014, Newland et al., 2018a). A recommended configuration of the NSGA-II algorithm was used (Newland et al., 2018a), including a probability of cross-over of 0.9 and a probability of mutation of 0.0028, the latter equalling the inverse of the number of decision variables. The stopping criterion adopted is based on reaching a set number of generations, 310, which was the number that could be completed in 48 hours of runtime using a parallelised version of NSGA-II (see Newland et al., 2018a). To account for stochasticity in the Metronamica model during the optimisation process, the map comparison metrics are averaged across ten stochastically generated replicates of simulated land-use maps for each iteration of NSGA-II, as suggested by Newland et al. (2018a).

The parallelised version of NSGA-II was run on the Phoenix high performance computing resources operated by Research Services at the University of Adelaide. The optimisation runs are evenly distributed across a number of CPU processing cores, with a number of slave processes evaluating the objective functions of different members of the population, and a master process co-ordinating the search, using the NSGA-II genetic operators (selection, cross-over, mutation, non-dominated sorting and crowding distance) to evolve the population of solutions.

#### 4.3.1.4 Assessment

The utility of the proposed calibration method is evaluated by considering its objective performance, the quality of the simulated output and the realism of the calibrated parameters obtained, to determine a final, calibrated model. As part of the first step of this process, the Pareto front of solutions for the calibration and validation periods are evaluated against benchmark models, from which reference metrics are calculated, to determine if the output solutions have correctly captured modelled processes. Two benchmark models are used, the growing clusters neutral model (Van Vliet et al., 2013b), and the random constraint match neutral model (Hagen-Zanker and Lajoie, 2008). The benchmark metrics are calculated for FK and AWCE, as FKS has an implicit baseline at a value of zero, as summarised in Table 4.2.

Table 4.2. Benchmark metric values for different benchmark models

Period	Growing Clusters		Random Constraint Match	
	FK	AWCE	FK	AWCE
Calibration	0.884	0.035	0.895	0.058
Validation	0.924	0.014	0.923	0.025

Next, in order to enable the realism of output maps and calibrated model parameters to be assessed, five solutions are extracted from those on the Pareto front for further consideration. To capture the diversity of the performance based on the objectives, three solutions are selected that correspond to the best objective performance for each metric (i.e. the solution with the highest FK, the solution with the highest FKS, and the solution with the lowest AWCE). Two intermediate solutions are also evaluated, one balanced between FK and FKS that had poor AWCE performance, and one balanced between AWCE and FK with poor FKS performance.

The degree to which the simulated output maps have captured the observed processes of land-use changes is assessed by visual inspection of the simulated output maps for these five solutions compared with the corresponding data. In order to evaluate the realism of the resultant

parameters, key parameters that differentiate the solutions based on which are more consistent with process knowledge, including inertia points, conversion points for conversions to the class *residential*, conversion points for conversions to the class *permanent crops*, influence tails for influence on the class *residential*, and the accessibility parameters, are considered, as it is not possible to present this analysis for all 352 parameters.

The utility of the proposed method is also evaluated by comparing the results from the *seeded* approach introduced in this paper with those obtained using the *process-specific* solutions (Section 4.3.1.1) and the standard, *unseeded* implementation of NSGA-II for optimisation-based calibration, starting with a randomly generated population (see Newland et al., 2018a). The unseeded method is configured the same as the seeded, as described in Section 4.3.1.3 (same probability of cross-over, probability of mutation, population size), and run for the same number of generations.

In order to enable the performance of two approaches that produce a Pareto front of solutions (i.e. the seeded and unseeded optimisation approaches) to be compared in an objective fashion, the hyper-volume metrics (Zitzler, 1999) of the final Pareto fronts obtained are compared. The hyper-volume calculates the volume of the multi-dimensional region enclosed by the Pareto front and a reference point, which provides a quantitative measure of solution quality and diversity. For this research, the hyper-volume is calculated by scaling the objectives to the range [0, 1] based on the observed maximum and minimum values of the two Pareto fronts, to ensure a balanced comparison between the different objectives. The hyper-volume is calculated from the *nadir* point in the objective space, the worst performance expected for the metrics used. As the ranges are normalised, the nadir point is 0 for FK and FKS (as these objectives are maximised) and 1 for AWCE (as this objective is minimised).

The performance of the seeded and unseeded optimisation approaches is also evaluated based on the degree to which the calibrated models align with process knowledge using the approach outlined above (i.e. the realism of the resulting maps and selected parameters are assessed for five calibrated models corresponding to those with the highest FKS, highest FK and lowest AWCE, as well as those for two intermediate solutions). Finally, the computational efficiency of the two optimisation methods is compared based on the number of generations required before all the Pareto front solutions are considered plausible, defined by solutions outperforming the calibration benchmarks, and the hyper-volume metric calculated for certain



generations during the optimisation. At the conclusion of the evaluation, a final, calibrated model is recommended.


#### **4.4 Results and discussion**

The results presented are divided as follows. First, the utility of the proposed seeded optimisation approach is evaluated by considering the objective performance metrics for the calibration and validation periods, the visual assessment of the simulated output maps for the five selected solutions, and an evaluation of a key set of resultant parameters that differentiate solution quality. Following this, the performance is compared with the two alternative calibration methods, unseeded optimisation and process-specific, considering objective performance, simulated output, parameter validity, and efficiency of the optimisation process. Following the evaluation a final, calibrated model is recommended.

##### ***4.4.1 Performance of seeded approach***

###### ***4.4.1.1 Objective performance***

The three-dimensional objective performance of the seeded approach for the calibration period is shown in Figure 4.5 (two-dimensional cross-section plots of the objectives are shown in comparison plots with the other two methods in Figure 4.13). In total, 239 Pareto optimal solutions were identified, with the shading of the blue indicating the relative performance for AWCE (lighter points correspond to lower (better) AWCE values). The black lines indicate the performance benchmark metric values, with the solid and dashed lines showing the growing clusters and random constraint match benchmarks, respectively. In Figure 4.5, the benchmarks for AWCE are not shown, as the errors obtained from the benchmark models are much greater than the values obtained for the seeded approach in the calibration period (see Table 4.2). As shown, a meaningful trade-off was found between the objectives, particularly for FK vs AWCE and FKS vs AWCE. All Pareto optimal solutions outperformed the benchmark metrics for the calibration period.



*Figure 4.5. Objective performance of the output from the seeded method for the calibration period. Solutions analysed in more detail are highlighted in black. The degree of shading of the dots indicates the AWCE performance, with lighter dots corresponding to lower (better) AWCE values. The solid and dashed lines correspond to the growing clusters and random constraint match performance benchmark, respectively (benchmark error values for AWCE are not shown as they were significantly higher than the plot scale).*

Figure 4.6 shows the performance of the seeded approach for the validation period with a three-dimensional plot of the objective performance metrics (two-dimensional cross-section plots of the objectives are shown in Figure 4.14). In Figure 4.6, the benchmark plane for AWCE is included for the growing clusters model. As shown, there is less spread in the solutions for the validation period, though the performance is still generally good, as a majority (86%) of the solutions outperform the performance benchmarks. The solutions that did not meet the performance benchmarks for validation were due to having higher AWCE values than the growing clusters benchmark model.





Figure 4.6. Objective performance of the output from the seeded method for the validation period. The degree of shading of the dots indicates the AWCE performance, with lighter dots corresponding to lower (better) AWCE values. The solid and dashed lines correspond to the growing clusters and random constraint match performance benchmark, respectively.


#### 4.4.1.2 Simulated output evaluation

As mentioned in Section 4.3.1.4, five solutions were analysed, shown by the black dots in Figure 4.5 and henceforth referred to as Solution 1 to Solution 5, with Solution 1 corresponding to the point shown in Figure 4.5 with the best FKS and FK performance, but the highest AWCE, to Solution 5 corresponding to the point shown in Figure 4.5 with the lowest AWCE, but relatively poor performance in terms of FKS and FK. These solutions were analyzed in further detail by considering the simulated output maps generated for the calibration period. These are shown in Figure 4.7, with 4.7(a-1) showing the data map for the end of the calibration period (2000), and the remaining maps showing simulated output for each solution, ordered by decreasing FKS value. Also shown are agreement maps between the data and the simulated output for the classes *residential* (red in the land-use map) and *industry & commerce* (purple in the land-use map), as these are the major urban classes and capturing their behavior accurately is a major focus of calibration.


From the simulated output evaluation, Solution 1 (best FKS) and Solution 2 (Intermediate 1), with corresponding simulated output maps 4.7(b) and 4.7(c) are most consistent with expectation, given the more accurate replication of the data, particularly the highlighted urban classes. This is discussed further for the classes *residential*, *industry & commerce*, and *airports*.




(a-1)




(b-1)




(b-2)




(b-3)




(c-1)




(c-2)




(c-3)



(d-1)

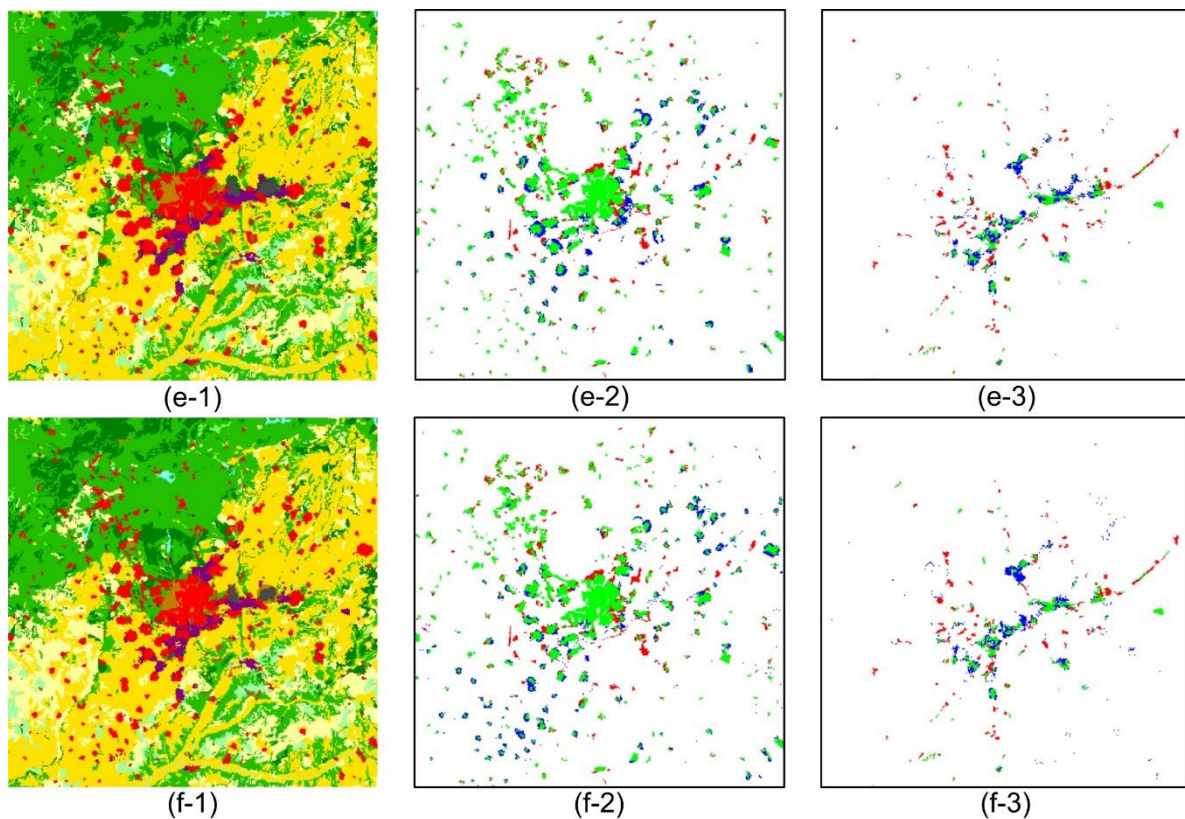


(d-2)



(d-3)





### Legend

Land-use map			Agreement map				
	Natural areas		Residential		Seaports		In neither map
	Arable land		Industry & commerce		Airports		In both maps
	Permanent crops		Recreation areas		Mine & dump sites		Only in data map
	Pastures		Forest		Fresh water		Only in simulated map
	Other agriculture		Road & rail				

Figure 4.7. Simulated output maps for the calibration period for five selected solutions for the seeded method. (a-1) is the data map for 2000. The remaining maps are ordered into rows by decreasing FKS value, from (b) the output for Solution 1 (best FKS), to (f) the output for Solution 5 (best AWCE). The first column (e.g. (b-1)) shows the simulated output map for the selected solution, the second column (e.g. (b-2)) shows the agreement of the class residential between the data and corresponding simulated output, and the third column (e.g. (b-3)) shows the agreement of the class industry & commerce between the data and corresponding simulated output

As shown in Figure 4.7, the simulated output across the solutions analysed exhibits a general tendency for the *residential* areas to be more clustered in the simulated output than in the data, and tends to favour the existing larger *residential* clusters, suggesting the solutions place an over-emphasis on the attraction of *residential* to itself. In addition, some of the solutions have not captured the correct transitions for the allocation of new *residential* areas. This is most noticeable in 4.7(f-2) and 4.7(e-2). In these solutions, the conversion from *arable land* to *residential* appears to be more common than in the data, while 4.7(b-2) and 4.7(c-2) appear the

most consistent with the data, having captured the conversions to *residential* from *arable land* and *natural areas*, which is consistent with these solutions corresponding to higher FKS values.

The appearance of the *industry & commerce* (purple) areas in the simulated output is also generally more clustered than in the data. This clustering is less prominent in Figures 4.7(b-3) and 4.7(c-3), whereas Figures 4.7(e-3) and 4.7(f-3) show much larger clusters of *industry & commerce*, particularly in the northern area of the main urban region, which, as shown in Figure 4.7(f-1), is near an area of *residential*, which would not be expected. When analysing the historic data, new areas of *industry & commerce* tended to be located near major roads. Some of the solutions did not appear to capture this attraction, and the ones that did, 4.7(f-3), and to a lesser extent 4.7(e-3), seem to favour an allocation near different roads than expected based on the data. The latter issue could potentially be resolved by further differentiating the various road types, as only one type of road was included.

Figure 4.7 also shows behaviour for the *airports* (grey) in the eastern portion of the land-use map that is not consistent with expectation. Airports tend to be near urban areas, but are not expected to be attractors of *residential* or *recreation areas*, or attractors of natural classes, such as *forest*. However, there appears to be an attractive influence for the different classes for certain solutions to the nearby neighbourhood of the *airport*. Figures 4.7(b-1), 4.7(c-1) and 4.7(d-1) show the airport being surrounded by *forest* (dark green), while Figure 4.7(f-1) shows that *recreation* (brown) completely surrounds the *airport* area. 4.7(e-1) shows the *airport* being surrounded by *industry & commerce* and *forest*, and while the former land-use is known to be attracted to airports, this behaviour was not found in the historic data. The obtained parameters in the discussed solutions confirmed the attraction of the different land-uses to the *airport*, so a solution to eliminate this behaviour in future would be to exclude such unrealistic interactions from the parameter space during the optimisation (preserving the parameter elimination from the process-specific approach). Additionally, introducing metrics such as the enrichment factor (Van Vliet et al., 2013b), which focus on neighbourhood composition, as an optimisation objective could potentially improve the resultant output.

As detailed previously, a majority of the solutions exhibited overly large clusters of the urban land-use classes compared to the data, which suggests that landscape pattern structure has not been as effectively captured as possible to generate simulated output consistent with process understanding. This is particularly noticeable given that Solutions 4 and 5, the two solutions analysed with the lowest AWCE, had simulated output that was not as consistent with process

knowledge as the other solutions analysed. This suggests that using area-weighting to aggregate the class level clumpiness errors may not be the most effective strategy that could be used to generate simulated output consistent with process understanding, as this focussed too much on capturing the behaviour of the active classes that occupied a majority of the area in the land-use map (see Contingency table, Supplementary material 4B). Given the relatively small area occupied by the urban classes, they had limited influence on the resultant metrics during the calibration procedure. Hence, the simulated output could potentially be improved by adjusting the aggregation strategy used, for example, by using a weighting scheme with a greater emphasis on the urban classes. This would ensure a greater focus on capturing the error associated with these classes, potentially producing simulated output more consistent with process understanding. Also, additional metrics could be used during the automatic calibration procedure that emphasize different pattern aspects, such as the fractal dimension or edge density (McGarigal, 2014).

#### 4.4.1.3 *Parameter evaluation*

The inertia parameters obtained for the different seeded solutions were generally consistent with expectation, as shown in Figure 4.8. However, of those analysed, Solution 1 is the most consistent with process understanding. As shown, the general trend is for the urban classes (*residential* and *industry & commerce*) to exhibit the highest inertia, and agricultural land-use classes *arable land*, *pastures*, and *other agriculture*, to exhibit the lowest inertia, with the other classes exhibiting moderate inertia. While this general trend is observed across the solutions, Solution 1 has better captured this dynamic, as the class *pastures* does not exhibit high inertia, unlike the other cases.






Figure 4.8. Comparison of inertia point influence values per class for seeded solutions

The conversion point parameters obtained for the seeded method were another example of results being generally consistent with expectation, as shown in Figure 4.9, with relatively high conversion values from the classes *arable land* and *other agriculture*. However, of the solutions analysed, Solution 3 is the most consistent with expectation. This is because for the other solutions, certain conversions are higher than expected, such as conversion from the class *industry & commerce* for Solutions 1 and 2, and from *permanent crops* for Solution 5. These conversions potentially explain the more clustered *residential* areas observed in the simulated output maps (Figure 4.7).




Figure 4.9. Comparison of conversion point influence values per class for transitions to the class residential for the seeded solutions

Conversions to the class *permanent crops* are shown in Figure 4.10. Across the solutions analysed, these parameters were generally positive, which might be expected for conversions from the classes *natural areas* and *arable land* (which are the dominant conversions across the solutions) but not for the classes *residential*, *recreation areas*, and *forests*, as is the case for Solutions 2 and 5.




Figure 4.10. Comparison of conversion point influence values per class for transitions to the class permanent crops for the seeded solutions.

Influence tails between the different land-use classes were generally consistent with expectation, though this varied across the solutions evaluated. An example of this is shown for the cross-influence tails to the class *residential* for the solutions analysed, shown in Figure 4.11 for the influence of urban classes (*residential*, *industry & commerce*, *recreation areas*). As shown, Solution 5 exhibits much lower self-influence of the class *residential* than the other solutions analysed, such that the self-influence is less than the cross-influence of *industry & commerce*, which is not consistent with expectation.

The influence tails also show that the seeded optimisation did not always preserve the parameter elimination from the initial population of process-specific and sampled solutions with the implementation used. Certain conversions from *airports* and *mine & dump sites* to *residential* (Supplementary material 4A) are not consistent with expectation, but are seen in the final model parameterisation and explain certain formations in the simulated output (see Section 4.4.1.2).

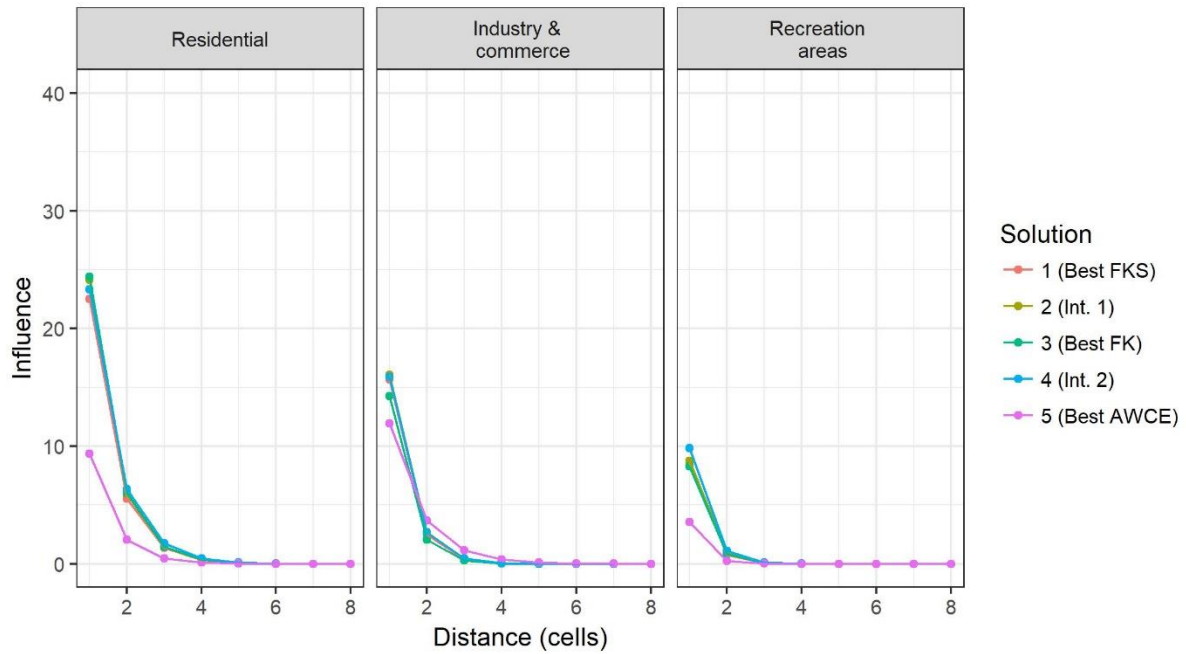


Figure 4.11. Comparison of neighbourhood influence tails for urban classes to the class residential for the seeded solutions

The accessibility parameters are shown in Figure 4.12 for the urban classes *residential*, *industry & commerce*, and *recreation areas*, with the weight and distance-decay parameters being similar across the classes and solutions, though the distance-decay for *recreation areas* was generally lower than for the other two. Though this is consistent with expectation, the parameter values are fairly equal, consistent with evaluation of the simulated output in that roads had a limited influence on the resultant output (Section 4.4.1.2).




Figure 4.12. Comparison of accessibility parameter values for classes residential, industry & commerce, and recreation areas for the seeded solutions

Given that certain parameters were calibrated to values inconsistent with process understanding suggests that such parameters, for example, the conversions from *permanent crops* to *industry & commerce* and the influence of *mine & dump sites* on *residential* had a limited influence on the simulated output and hence the objectives of model performance for the calibration period, meaning they were not emphasized during calibration. However, parameters that are inconsistent with process understanding limit model validity, and the potential for application to long-term scenario analysis. Hence, further improvements could be made to the automatic calibration method to improve model validity. An example would be to place more constraints on certain parameters during the optimisation, which could be based on the parameter elimination from the process-specific approach, to limit the possible ranges of the parameters. Alternatively, these parameters could be eliminated entirely from automatic calibration, which could also make the optimisation more efficient.


#### 4.4.2 Comparative assessment

##### 4.4.2.1 Objective performance


A comparison of the objective performance of the different calibration methods is presented in Figure 4.13 for the calibration period, comparing the 239 Pareto optimal points for the seeded method (blue dots) with the 306 Pareto optimal points identified using the unseeded method (red dots) and the two process-specific solutions (green dots). Figure 4.13(a) shows the three-

dimensional objective space and indicates the benchmark metric values as planes in the three-dimensional space. Figures 4.13(b)–4.13(d) show a two-dimensional comparison of the trade-off between the different combinations of performance objectives, FKS against AWCE (b), FK against AWCE (c) and FK against FKS (d). Figures 4.13(b)–4.13(d) indicate the solutions that were subsequently evaluated, with black dots for the seeded approach and black diamonds for the unseeded approach.


Figure 4.13 shows that the seeded optimisation method outperformed both other calibration methods in the objective space for the calibration period. The starting position of the process-specific solutions has clearly been improved upon, and the Pareto front has better objective performance and more diversity than the Pareto front of the unseeded approach. This is particularly observed in Figure 4.13(c) and 4.13(d) where, as seen in the bi-objective comparison, the Pareto solutions from the proposed method completely dominate all Pareto solutions from the unseeded approach. The resultant hyper-volume value reinforces this (scaled to the minimum and maximum values observed in Figure 4.13), with a value of 0.717 for the seeded approach, versus 0.358 for the unseeded.




(a)



(b)



(c)




(d)


Figure 4.13. Comparison of objective performance for the calibration data for the optimisation methods run for 310 generations, and the process specific solutions. (a) a 3-D plot of the objective space; (b) a cross-section of the bi-objective space for FK vs AWCE; (c) a cross section of the bi-objective space for FKS vs AWCE; and (d) a cross section of the bi-objective space for FKS vs FK. Solid and dashed lines indicate the growing clusters and random constraint match neutral model benchmarks respectively. Black dots and diamonds show analysed solutions for the seeded and unseeded methods respectively.

The improved objective performance of the seeded solutions over the other two methods was even more pronounced for the validation period (shown in Figure 4.14), suggesting that the results for the seeded calibration method better generalise to the case study. As shown in all bi-objective comparisons, the Pareto solutions from the proposed method dominate those from the unseeded method with regard to FK and FKS, with limited deterioration in the AWCE values. The hyper-volume metric reinforces this, with a value for the seeded method (with scaling to the minimum and maximum values observed) of 0.728, versus 0.490 for the unseeded method.







(a)



(b)



(c)




(d)


Figure 4.14. Comparison of objective performance for the validation data for the optimisation methods run for 310 generations and the process specific method: (a) a 3d plot of the objective space; (b) a cross-section of the bi-objective space for FK vs AWCE; (c) a cross section of the bi-objective space for FKS vs AWCE; and (d) a cross section of the bi-objective space for FKS vs FK. Solid and dashed lines indicate the growing clusters and random constraint match neutral model benchmark respectively

#### 4.4.2.2 *Simulated output comparison*


Comparison of the simulated output maps via visual interpretation suggests that the seeded approach was more consistent with the data than the unseeded and process-specific solutions. The simulated output maps are presented in Figure 4.15 for the selected solutions for the unseeded method, and Figure 4.16 for the process-specific solutions. The seeded output performs better because the clustering of the urban land-uses *residential* (red) and *industry & commerce* (purple) is more consistent with the data than for the other methods. Of the alternative methods, the unseeded output (Figure 4.15) generated simulated maps with more clustering of the major urban classes *residential* (red) and *industry & commerce* (purple) across the solutions analysed than the data and the corresponding seeded output. For the process-specific solutions (Figure 4.16), the solution with a slight preference for locational agreement exhibited less clustering of the urban classes analysed (though still more than the data) but removed *residential* area in the south-west region, shown in Figure 4.16(b-2), which is not consistent with expectation. The solution with a slight preference for landscape pattern structure exhibited a similar issue, and also resulted in large clusters of *industry & commerce* in totally different areas than observed in the data, shown in Figure 4.16(c-3), though this solution did appear to include sufficient attraction to the relevant roads in the allocation of *industry & commerce*.




(a-1)




(b-1)




(b-2)




(b-3)




(c-1)




(c-2)




(c-3)



(d-1)

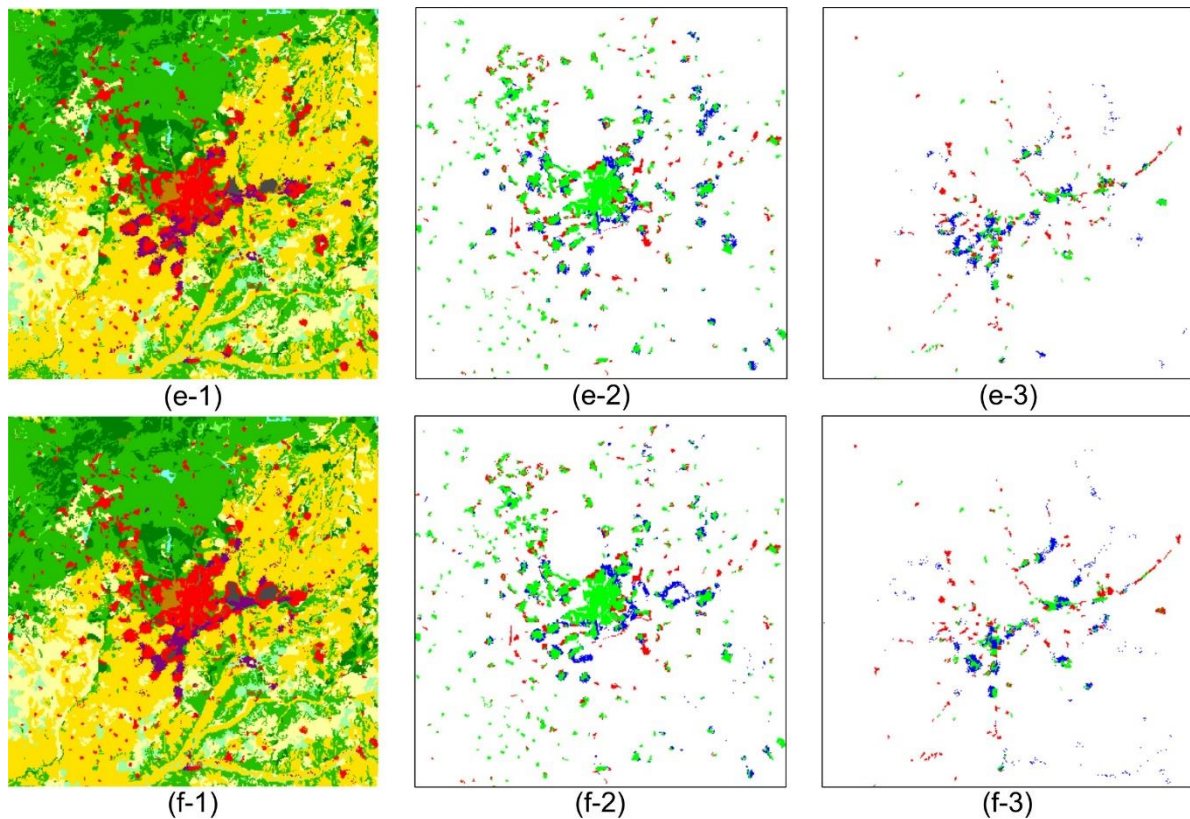


(d-2)



(d-3)

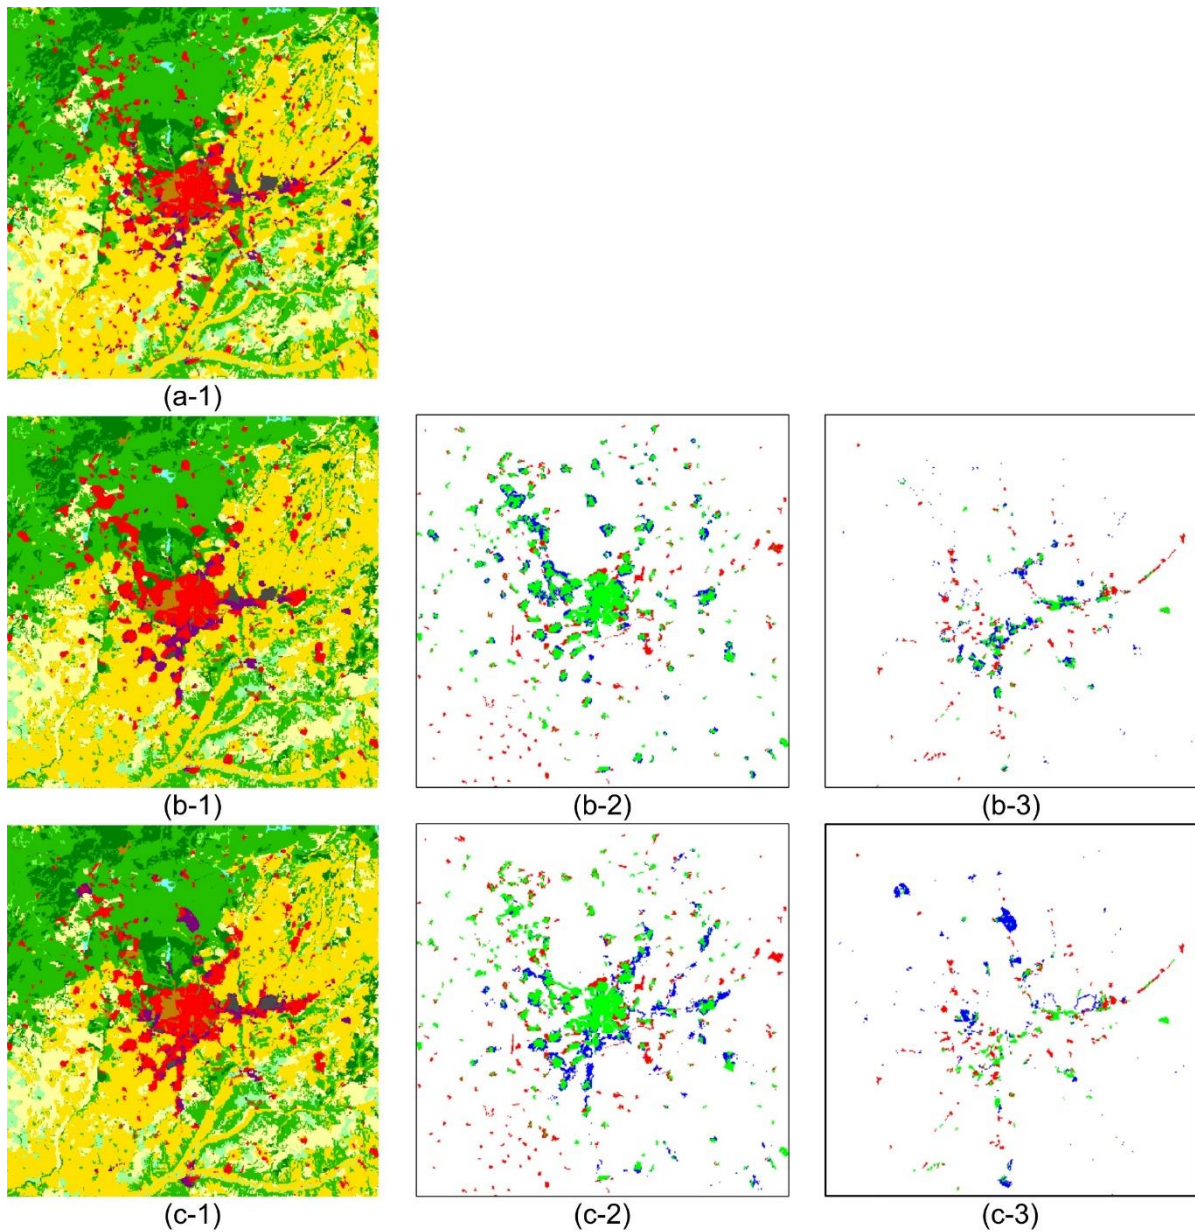




### Legend

Land-use map			Agreement map				
	Natural areas		Residential		Seaports		In neither map
	Arable land		Industry & commerce		Airports		In both maps
	Permanent crops		Recreation areas		Mine & dump sites		Only in data map
	Pastures		Forest		Fresh water		Only in simulated map
	Other agriculture		Road & rail				

Figure 4.15. Simulated output maps for the calibration period for five selected solutions for the unseeded method. (a-1) in the data map for 2000. The remaining maps are ordered into rows by decreasing FKS value. The first column (e.g. (b-1)) shows the simulated output map for the selected solution, the second column (e.g. (b-2)) shows the agreement of the class residential between the data and corresponding simulated output, and the third column (e.g. (b-3)) shows the agreement of the class industry & commerce between the data and corresponding simulated output



### Legend


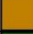

Land-use map			Agreement map				
	Natural areas		Residential		Seaports		In neither map
	Arable land		Industry & commerce		Airports		In both maps
	Permanent crops		Recreation areas		Mine & dump sites		Only in data map
	Pastures		Forest		Fresh water		Only in simulated map
	Other agriculture		Road & rail				

Figure 4.16. Simulated output maps for the process-specific calibration method compared with the calibration data (a-1), that resulted in better locational agreement performance (b-1) to (b-3) and better landscape pattern structure performance (c-1) to (c-3) for the calibration period. Maps in the middle column (e.g. (b-2)) show the agreement of the class residential between the simulated output and data. Maps in the right-most column (e.g. (b-3)) show the agreement of the class industry & commerce between the simulated output and data.

#### 4.4.2.3 Parameter comparison

The calibrated parameters obtained from the seeded method were more consistent with expectation than those obtained for the unseeded method and similar with the process-specific parameters. Hence, comparisons are only presented between the seeded and unseeded method. This is illustrated for the parameters previously discussed (Section 4.4.1.3), with comparison plots presented in Supplementary material 4A.

Generally, the calibrated parameter values obtained using the unseeded method were higher than those of the corresponding seeded parameters, leading to parameter values that are less consistent with expectation. This was exhibited for the inertia point parameters for the unseeded solutions, with agricultural classes (such as *pastures* and *other agriculture*) exhibiting inertia values that were approximately equivalent to those of the urban classes (*residential, industry & commerce, recreation areas*), which is inconsistent with expectation, as agricultural land-uses generally exhibit less inertia to facilitate allocation of new urban classes. Hence, such parameter values could limit long-term application. This led to other conversion and influence parameters being much higher than expected, which potentially explains why more clustering was observed in the simulated output (Figure 4.15). Additionally, the unseeded method generated more parameter values across solutions that were less consistent with process understanding, such as a high conversion parameter for *residential* to *permanent crops* and vice-versa, and higher attraction between infrastructure classes and urban classes (highlighted in Supplementary material 4A for the influence of *airports* on *residential* and *road & rail* on *residential*).

#### 4.4.2.4 Optimisation efficiency

The optimisation process was more efficient for the seeded approach, when considering the computational effort to generate a full Pareto front of solutions considered valid for the calibration period with respect to the benchmark tests, and the hyper-volume metric (scaled to the minimum and maximum values observed in Figure 4.13). This is shown in Table 4.3, which shows the percentage of solutions in the Pareto front for a set of generations that outperform the benchmarks, and hence would be considered valid. As shown, the seeded approach generated plausible models with far less computational effort, requiring only 20 generations for the Pareto front to consist entirely of plausible solutions. By contrast, over 80 generations are required for this to happen with the unseeded approach. The efficiency of the seeded approach is also highlighted by the hyper-volume metric, which at generation 80 for the seeded method is greater than the value at generation 310 for the unseeded method.

Table 4.3. Comparison of percentage of Pareto front solutions outperforming benchmark performance metrics and hyper-volume for seeded and unseeded calibration methods

Generation	Pareto front solutions outperforming performance benchmarks (%)		Hyper-volume	
	Seeded	Unseeded	Seeded	Unseeded
0	0%	0%	0.015	0.000
10	93%	0%	0.136	0.000
20	100%	0%	0.217	0.000
40	100%	78%	0.331	0.064
80	100%	99%	0.447	0.180
160	100%	100%	0.550	0.283
310	100%	100%	0.717	0.358

#### 4.4.3 Determination of final calibrated model

Of the (semi) automatic calibration methods compared in this analysis, the seeded method performed better than the unseeded and process-specific method. The objective performance is superior to the other methods, especially for the validation period, suggesting better generalised behaviour, and the corresponding simulated output and parameter values reflect this, as they were more consistent with process understanding than the other two methods. Of the seeded solutions evaluated, the best solution corresponded to Solution 1, generating the simulated output shown in Figure 4.7(b), which was most consistent with the data, and had parameters that tended to be more consistent with expectation across the evaluated seeded (and other) solutions. The recommended parameterisation for the case study is given in Supplementary material 4C.

#### 4.5 Conclusions

To address issues with the computational efficiency and parameter validity of optimisation-based approaches applied to automatically calibrate transition potential based LUCA models, this research proposed a method to integrate process understanding into optimisation-based automatic calibration. The proposed method was used to generate an initial population of solutions, based on the output of a process-specific semi-automatic calibration method, to initialise a multi-objective optimisation algorithm with a diverse set of solutions, some of which occurred in good regions of the parameter space. The presented method is generic, allowing for the use of different optimisation algorithms, land-use models, and process-specific calibration methods.



The performance of the seeded approach was demonstrated via application to a case-study of Madrid, Spain, with its utility compared with that of a standard, unseeded implementation of an optimisation-based automatic calibration method, and a process-specific semi-automatic calibration method. The key advantage of the seeded approach over the other two methods was superior objective performance, identifying more diverse Pareto front solutions with better performance (as quantified by the objectives) for both the calibration and validation data. This translated to the seeded method producing simulated output maps and calibrated parameters that were more consistent with process understanding. The seeded approach also had an efficiency advantage in the generation of plausible calibrated models with respect to the benchmarks, requiring less computational effort to generate an entire Pareto front of plausible solutions.

Given the advantages achieved using a seeded approach, further avenues utilising such an approach should continue to be explored. Further improvement could be achieved by integrating greater parameter reduction in the seeded optimisation approach, by preserving the parameter elimination from the process-specific method used for this specific implementation, or based on discursive knowledge, to reduce the parameters included for optimisation-based calibration. Additionally, further improvements could be achieved by using different metrics as calibration objectives during optimisation that evaluate different aspects of the land-use map. These improvements could lead to more plausible model parameterisations being generated, and improve the efficiency of the method. These advantages presented have the potential to make the application of optimisation-based automatic calibration methods for LUCA models more efficient and effective in the support of policy development for long term planning.

#### **4.6 Acknowledgements**

The authors wish to acknowledge the financial support from the Bushfire and Natural Hazards Cooperative Research Centre, made available by the Commonwealth of Australia through the Cooperative Research Centre program. This work was supported with supercomputing resources provided by the Phoenix High Performance Computing services at the University of Adelaide.

## 4.7 Supplementary material

### 4.7.1 Supplementary material 4A: Comparison plots of parameters analysed for seeded and unseeded calibration method

This section contains a series of plots comparing the parameters obtained from the seeded and unseeded calibration methods, the inertia point parameters, the conversion point parameters for transitions to the class *residential*, conversion point parameters for transitions to the class *permanent crops*, influence tails for interactions for conversions to the class *residential*, and the accessibility parameters. Generally, the unseeded calibration method resulted in higher parameter values. The unseeded method also resulted in more examples of parameters that were not consistent with expectation, for example, high inertia across a majority of classes (Figure 4A.1), high conversion values for transitions from the classes *forest* and *permanent crops* to *residential* (Figure 4A.2), greater emphasis on transitions to *permanent crops* from the class *residential* (Figure 4A.3), attractive influences on *residential* land-use exerted by classes such as *airports* and *seaports* (Figure 4A.7), and no weighting given to *residential* land-use by accessibility (Figure 4A.8).




Figure 4A.1. Comparison of inertia influence per class for analysed solutions for seeded and unseeded calibration methods




Figure 4A.2. Comparison of conversion point influence values to class residential for each land-use class for analysed solutions for seeded and unseeded calibration methods




Figure 4A.3. Comparison of conversion point influence values to class permanent crops for each land-use class for analysed solutions for seeded and unseeded calibration methods

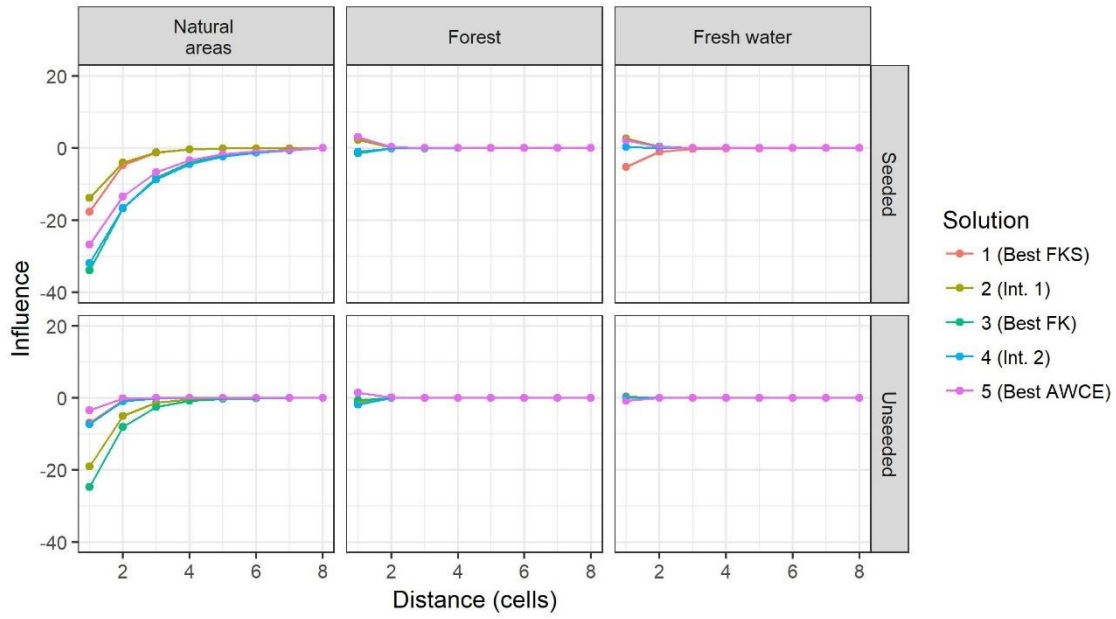


Figure 4A.4. Comparison of tail influences to the class residential for the natural land-use classes for analysed solutions for seeded and unseeded calibration methods




Figure 4A.5. Comparison of tail influences to the class residential for the agricultural land-use classes for analysed solutions for seeded and unseeded calibration methods

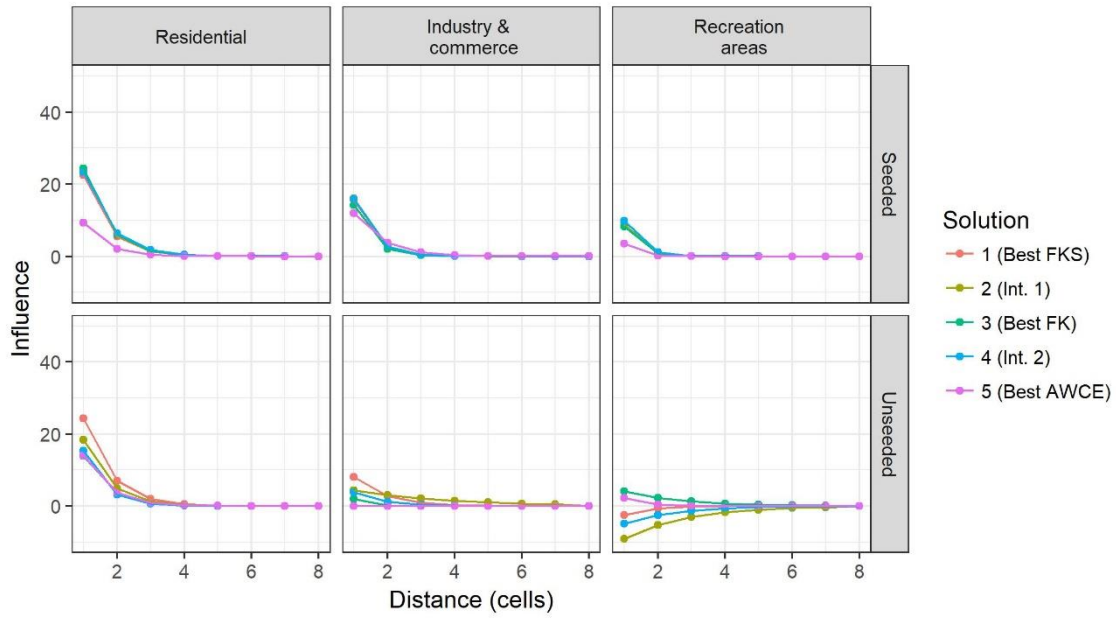


Figure 4A.6. Comparison of tail influences to the class residential for the urban land-use classes for analysed solutions for seeded and unseeded calibration methods




Figure 4A.7. Comparison of tail influences to the class residential for the infrastructure land-use classes for analysed solutions for seeded and unseeded calibration methods

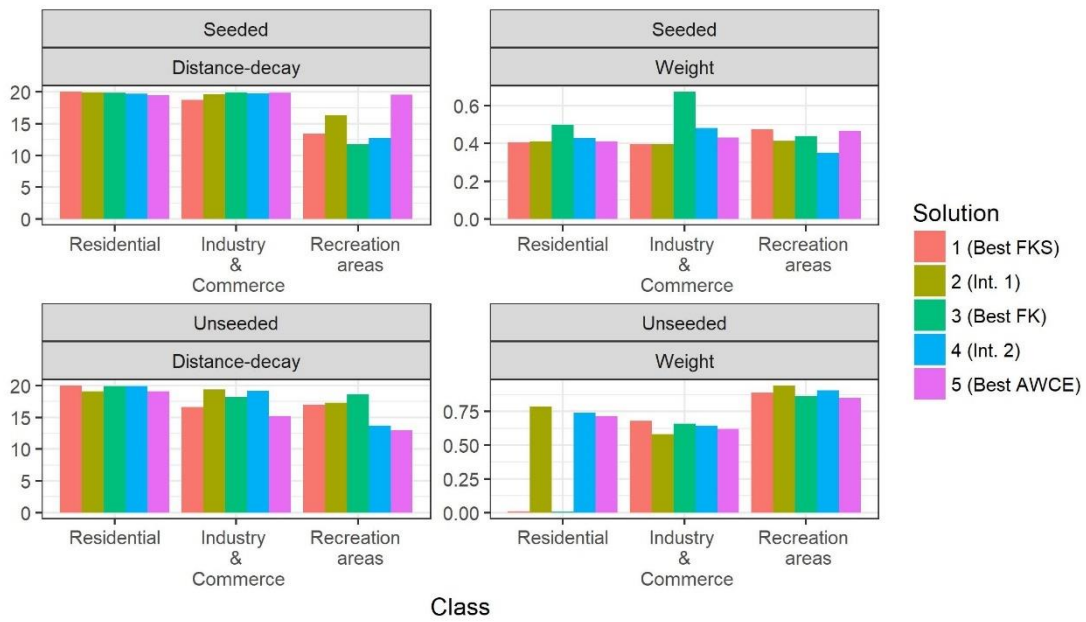


Figure 4A.8. Comparison of accessibility parameters for the urban land-use classes for analysed solutions for seeded and unseeded calibration methods

#### 4.7.2 Supplementary material 4B: Contingency table

This section contains a contingency table for the data maps for the Madrid case study for the calibration (1990-2000) period.

Table 4B.1. Madrid contingency table, 1990-2000

		Map 2000															
	LUC	NAT	ARL	PER	PAS	OAG	RES	I&C	REC	FOR	R&R	POR	AIR	M&D	FRE	MAR	TOT
Map 1990	NAT	46821	268	3	0	50	1051	316	164	357	31	0	12	161	20	0	49254
	ARL	1540	55265	30	0	383	1492	1059	121	32	70	0	92	351	26	0	60461
	PER	50	18	4278	0	0	29	3	0	0	0	0	0	18	5	0	4401
	PAS	13	0	0	624	0	23	5	0	0	0	0	0	0	0	0	665
	OAG	191	34	15	0	20761	207	45	22	6	5	0	0	101	0	0	21387
	RES	1	4	0	0	6	7653	85	31	0	1	0	0	0	0	0	7781
	I&C	3	8	0	0	0	15	1152	5	0	0	0	0	0	0	0	1183
	REC	0	0	0	0	0	37	0	656	0	0	0	0	0	0	0	693
	FOR	55	16	0	0	8	21	0	11	12164	0	0	0	0	7	0	12282
	R&R	0	0	0	0	0	0	0	0	0	145	0	0	0	0	0	145
	POR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	AIR	0	0	0	0	0	0	4	0	0	0	0	494	0	0	0	498
	M&D	33	6	0	0	0	29	26	39	0	0	0	0	368	2	0	503
	FRE	24	4	0	0	0	0	0	0	0	0	0	0	0	719	0	747
	MAR	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
	TOT	48731	55623	4326	624	21208	10557	2695	1049	12559	252	0	598	999	779	0	160000

### 4.7.3 Supplementary material 4C: Final model parameterisation

This section contains the final, calibrated parameters for the Madrid case study

Table 4C.1. Inertia/Conversion parameters

From/To	ARL	PER	PAS	OAG	RES	I&C	REC	FOR
NAT	2.68	47.03	0.55	28.67	16.98	0.23	7.60	47.99
ARL	249.16	33.99	5.47	41.23	36.12	53.09	40.42	82.27
PER	5.13	852.27	6.38	7.27	5.61	7.73	33.68	2.47
PAS	10.02	1.34	300.84	24.44	32.77	0.06	12.41	35.97
OAG	49.07	23.07	7.43	304.17	51.45	43.53	74.89	15.39
RES	30.25	5.04	0.56	7.06	1350.05	4.17	66.99	18.98
I&C	1.07	3.46	15.83	2.08	35.93	992.73	6.37	1.61
REC	5.20	0.27	19.74	1.38	13.23	12.29	701.94	18.95
FOR	12.41	3.26	12.83	23.28	7.25	17.07	4.27	880.82
R&R	1.96	10.64	6.75	6.63	0.01	4.55	17.07	8.74
POR	13.21	10.35	0.56	27.60	14.63	20.38	24.04	6.66
AIR	1.82	1.20	13.10	42.78	12.72	29.52	0.06	4.82
M&D	10.94	29.87	4.24	11.89	12.52	2.51	58.30	8.42
FRE	0.01	12.06	3.63	25.85	29.76	17.68	0.29	3.19

Table 4C.2. Neighbourhood rule a parameters

From/To	ARL	PER	PAS	OAG	RES	I&C	REC	FOR
NAT	-31.91	1.41	7.92	-20.93	-65.35	-4.79	5.86	-3.99
ARL	34.84	28.40	-30.06	6.63	34.39	19.79	-19.00	5.07
PER	2.63	61.31	-42.36	27.61	6.94	-3.23	-2.61	-21.57
PAS	9.17	-8.67	86.80	13.78	94.37	16.27	-7.61	27.16
OAG	6.82	44.57	-4.01	76.12	-19.17	-10.54	-1.90	18.37
RES	-76.40	28.58	-3.62	-66.99	91.60	12.41	23.32	-29.36
I&C	-13.64	44.29	14.05	6.79	98.63	43.51	19.69	-4.16
REC	-13.83	-19.45	-18.26	-39.72	98.83	12.69	98.19	-3.29
FOR	23.85	-40.39	-0.48	-28.76	38.92	-30.36	-73.53	57.09
R&R	-27.21	6.79	-31.67	55.99	8.66	17.71	10.73	-24.12
POR	-51.59	2.52	-48.00	-11.05	26.88	-31.46	18.74	1.79
AIR	-20.08	1.68	-9.39	1.38	1.60	44.64	5.18	69.36
M&D	18.44	81.39	21.24	-3.07	28.09	77.74	82.08	8.99
FRE	11.73	-1.30	-10.35	-3.78	-26.25	18.32	20.19	-6.57



Table 4C.3. Neighbourhood rule *b* parameters

From/To	ARL	PER	PAS	OAG	RES	I&C	REC	FOR
NAT	1.00	2.69	3.32	2.76	1.31	3.81	1.72	3.42
ARL	3.38	3.17	1.18	2.81	4.18	2.29	3.22	1.44
PER	2.16	1.77	2.38	2.22	2.71	2.77	1.79	2.27
PAS	3.30	2.02	2.21	3.01	2.79	2.60	3.72	2.79
OAG	2.55	2.35	1.77	2.27	0.65	1.14	3.19	2.59
RES	1.28	1.47	2.13	1.14	1.40	0.90	1.53	4.01
I&C	2.45	2.97	2.62	2.17	1.84	0.81	2.20	3.84
REC	1.92	2.21	4.27	2.65	2.44	0.73	0.84	1.00
FOR	3.02	0.80	2.37	2.64	2.59	3.01	2.16	1.72
R&R	0.20	1.90	2.67	2.59	2.42	0.68	2.22	2.72
POR	1.92	2.80	2.01	1.28	3.04	2.51	3.98	3.00
AIR	2.35	2.30	2.41	2.27	3.97	1.56	0.50	0.70
M&D	2.53	1.36	2.15	1.41	1.12	2.01	0.86	1.57
FRE	3.20	2.72	2.02	1.30	1.63	2.95	2.28	1.77

Table 4C.4. Accessibility parameters

Land-use class	Major roads distance decay	Major roads weight
ARL	2.06	0.00
PER	2.72	0.00
PAS	1.39	0.00
OAG	0.10	0.00
RES	20.00	0.41
I&C	18.75	0.40
REC	13.42	0.47
FOR	1.81	0.00



## **5 Conclusions**

Given the value that Cellular Automata (CA) based land-use models can provide in the development of policy and spatial planning, and the availability of generic Land Use Cellular Automata (LUCA) models for direct application to the study region of interest, application of LUCA models has gained increased research focus recently. Of the three major types of LUCA model frameworks, pattern extrapolation, development probability and transition potential, transition potential models have gained increasing attention, especially in relation to their calibration because transition potential based LUCA models are traditionally calibrated manually, which is time-consuming, subjective, and difficult to repeat. Hence, there is a large focus on the development of automatic calibration methods for transition potential based LUCA models.

Automatic calibration requires objective measures of model performance, and for transition potential based LUCA models, two distinct aspects must be considered, locational agreement and landscape pattern structure. Hence, metrics quantifying both aspects should be included for automatic calibration. Additionally, parameter dimensionality, the number and range of possible parameters, plays a large role in the complexity of automatic calibration, as transition potential based LUCA models generally feature hundreds of parameters for calibration.

In order to address these issues, two main forms of (semi) automatic calibration methods have been developed for transition potential based LUCA models. The first are optimisation-based approaches, which generally use a population-based metaheuristic, where a population of solutions is generated (i.e. a number of different sets of LUCA model parameter values), which are adjusted based on some operators to improve the objective performance of the solutions over a number of iterations, and are effective at generating multiple possible model parameterisations. The second are process-specific approaches, targeted to the parameters underlying a specific process (commonly the neighbourhood rules capturing spatial dynamics), aiming to generate a calibrated model that is consistent with process knowledge efficiently. This thesis has presented new methodologies for both types of approaches, as well as a method for combining the two approaches, to facilitate the (semi) automatic calibration of transition potential based LUCA models to support policy development and spatial planning.

### **5.1 Research contribution**

The major contribution of this work has been the development of a set of robust, generic methods for the (semi) automatic calibration of transition potential LUCA models. The

proposed methods facilitate (semi) automatic calibration using either of the two common approaches, and also allow for a combined approach where the benefits of both are utilised. The utility of each method has been demonstrated via case study applications, and each show promising potential for future applications of LUCA models to support long term planning and policy development. Specifically, in meeting the primary objectives of this research that were presented in the Introduction, the following research contributions were made:

A generic multi-objective optimisation framework for automatic calibration of transition potential LUCA models was developed in Paper 1 (Chapter 2). The framework is generic, allowing for substitution of the various components (e.g. the LUCA model, metrics, and the optimisation algorithm), utilises multi-objective optimisation to allow for the exploration of trade-offs between the LUCA model performance objectives, appropriately considers the inherent stochasticity included in LUCA models, and facilitates increased computational efficiency through its implementation. The capability of the generic framework was illustrated with an application to the Randstad region of the Netherlands. The results indicated that the method was able to effectively generate multiple plausible model parameterisations that, following further assessment and refinement, could be used for long-term analysis.

A process-specific semi-automatic calibration method that integrates objective analysis with discursive input to facilitate efficient calibration of neighbourhood rules with a limited computational budget (i.e. achievable using a desktop PC) was developed in Paper 2 (Chapter 3). The method first reduces the complexity of the calibration problem and then calibrates the remaining neighbourhood rules in a computationally efficient manner. The utility of the proposed approach was demonstrated via application to four European case studies with varying physical characteristics and rates of growth. The results indicated that the method was capable of generating calibrated models outperforming the provided benchmarks, with model parameters and simulated output that were consistent with empirical expectation.

A generic framework for hybrid automatic calibration, which integrates domain knowledge into a multi-objective optimisation approach, was developed in Paper 3 (Chapter 4). The hybrid approach, tailored to population-based metaheuristics, seeds the optimisation algorithm with a diverse initial population, including solutions identified using a process-specific method that occur in good regions of the parameter space and are consistent with process knowledge. The performance of the seeded approach was demonstrated via a case-study application to Madrid, Spain, resulting in multiple plausible model parameterisations that were consistent with process

understanding. The utility of the hybrid method was also compared with that of the standard, unseeded implementation of an optimisation-based automatic calibration method introduced in Chapter 2, and the process-specific method introduced in Chapter 3. The seeded method outperformed the other two approaches in terms of objective performance and solution quality that was more consistent with process knowledge.

This research also met the set of secondary objectives that were presented in the Introduction, naturally aligned with the development of automatic calibration methods for LUCA models, making the following research contributions:

Across the three papers, the impact of aggregating class level landscape pattern structure metrics, for this research measured using the clumpiness error, to a single measure, on the automatic calibration output was explored. Paper 1 used the average to aggregate the class-level errors of the clumpiness metric, which was found to over-represent minor land-use classes, and suggested a different aggregation method was required for conversion to a single value. Subsequently, the area-weighted clumpiness error was used in the case study application of the process-specific approach (Paper 2) and the hybrid approach (Paper 3), and was found to better capture the overall performance of the model calibration, though potential improvements could still be made by weighting the error more towards the main classes that are the focus of calibration (i.e. urban land-use classes).

The three papers also investigated the impact of the selected locational agreement metric(s) on the resultant calibrated model. Locational agreement metrics generally focus on either measuring the agreement of the entire land-use map, quantified in this research using Fuzzy Kappa (FK), or the subset of cells that transitioned, quantified in this research using Fuzzy Kappa Simulation (FKS). Paper 1 identified a potential trade-off between the FK and FKS metrics. Consequently, both were included as objectives for the process-specific method in Paper 2, and it was found that, for this approach, best performance generally required balancing between the two. The trade-off between the two metrics was confirmed in Paper 3, where three calibration objectives were used. Hence, it is important to capture both elements of locational agreement when using a (semi) automatic calibration method.

Finally, the development of the (semi) automatic calibration methods included the integration of discursive knowledge into the application and analysis. Papers 2 and 3 demonstrated how discursive knowledge can facilitate more efficient calibration. All three papers also showed that integrating empirical knowledge into the evaluation, by considering the realism of the

simulated output and the plausibility of the parameters obtained, allowed for more robust final calibrated models to be obtained. In turn, using an automatic approach also provided insights into the model structure and its behaviour, to improve the understanding of the LUCA model itself, and the interplay of factors driving land use change.

## **5.2 Limitations**

There were certain limitations of this research, which are discussed below.

1. There are a number of metrics for the quantification of locational agreement and landscape pattern structure. This research only considered three metrics, selected based on preference, including Fuzzy Kappa, Fuzzy Kappa Simulation, and clumpiness. As automatic calibration methods are driven by the performance objectives used, different metrics have the potential to influence the results obtained.
2. Applications of multi-objective optimisation-based automatic calibration methods were implemented using the Non-dominated Sorting Genetic Algorithm II. However, there is a range of available optimisation algorithms that could be used that could potentially improve the solutions obtained or find them in a more computationally efficient manner.
3. The neighbourhood rules were parameterised using shape functions to reduce the parameter dimensionality during automatic calibration. However, this removed the potential for more complex neighbourhood relationships to be obtained, which could potentially improve model performance.
4. The proposed methods have been developed to calibrate for two data maps with a certain time interval between them. The time interval between the data maps has the potential to impact the results obtained, particularly for the process-specific method (Paper 2) that bases the elimination of parameters on the available data. The impact of the duration between the data could be further explored to determine the impact on automatic calibration.
5. Results outperformed benchmarks of common growth strategies, but did not fully capture behaviour that was consistent with process knowledge for certain case studies, due to limited integration of this knowledge, which could improve the validity of the output of (semi) automatic calibration methods.
6. The broader application of a LUCA model included as part of a decision support system generally involves stakeholder input in the calibration process. Stakeholder

involvement in the calibration process was not considered, though each calibration method allows for discursive interpretation. This could be explored in future work.

7. Only a small number of predominantly urban case studies was tested, which potentially limits the generality of the results.
8. Automatic calibration was limited to parameters underlying neighbourhood rules and accessibility, not considering suitability or zoning, which could improve the resultant automatic calibration output.

### **5.3 Future work**

Based on the above limitations, the following future work is proposed.

1. Future application of the proposed methods using different metrics to quantify locational agreement and landscape pattern structure. This could include the use of a greater number of objectives (a maximum of three was used in this research) to further investigate the resultant trade-offs, which may be found for different landscape pattern structure metrics. The impact on solution quality could be investigated via discursive interpretation to determine which set(s) of metrics result in the most realistic performance.
2. Development of metrics to capture aspects of the discursive evaluation not currently captured by locational agreement or landscape pattern structure metrics, which could be included as automatic calibration objectives.
3. Further application of different optimisation algorithms, via substitution into the proposed multi-objective optimisation framework, to determine if superior performance can be achieved. This includes the use of metaheuristics, such as particle swarm optimisation (Blecic et al., 2015), or the use of hyper-heuristics (Maier et al., 2015) that combine different optimisation algorithms.
4. Further exploration of alternative parameterisations of neighbourhood rules to permit more complex shapes to be generated during automatic calibration.
5. Further investigation of the impact the time interval has on the resultant calibration. For the process-specific method, this could include evaluating the time interval in conjunction with the thresholds used for parameter elimination.
6. Further research into understanding land-use dynamics and the incorporation of discursive knowledge into automatic calibration methods.

7. The use of the automatic calibration frameworks as part of a participatory calibration approach (Hewitt et al., 2014), where stakeholder feedback is integrated into the assessment of the automatic calibration output.
8. Application of automatic calibration methods to a larger number of case studies with a more diverse range of characteristics and land-use classes.
9. Application of automatic calibration that includes suitability and zoning parameters, and integrating process knowledge of these parameters into seeding.



## 6 References

- ALJOUFIE, M., BRUSSEL, M., ZUIDGEEST, M., VAN DELDEN, H. & VAN MAARSEVEEN, M. 2016. Integrated analysis of land-use and transport policy interventions. *Transportation Planning and Technology*, 39, 329-357.
- BARREDO, J. I., KASANKO, M., MCCORMICK, N. & LAVALLE, C. 2003. Modelling dynamic spatial processes: simulation of urban future scenarios through cellular automata. *Landscape and Urban Planning*, 64, 145-160.
- BATTY, M. & LONGLEY, P. 1994. *Fractal cities: a geometry of form and function*, Academic Press Professional, Inc.
- BERBEROĞLU, S., AKİN, A. & CLARKE, K. C. 2016. Cellular automata modeling approaches to forecast urban growth for adana, Turkey: A comparative approach. *Landscape and Urban Planning*, 153, 11-27.
- BI, W., DANDY, G. C. & MAIER, H. R. 2015. Improved genetic algorithm optimization of water distribution system design by incorporating domain knowledge. *Environmental Modelling & Software*, 69, 370-381.
- BI, W., DANDY, G. C. & MAIER, H. R. 2016a. Use of domain knowledge to increase the convergence rate of evolutionary algorithms for optimizing the cost and resilience of water distribution systems. *Journal of Water Resources Planning and Management*, 142.
- BI, W., MAIER, H. R. & DANDY, G. C. 2016b. Impact of Starting Position and Searching Mechanism on the Evolutionary Algorithm Convergence Rate. *Journal of Water Resources Planning and Management*, 142, 04016026.
- BLECIC, I., CECCHINI, A. & TRUNFIO, G. A. 2015. How much past to see the future: a computational study in calibrating urban cellular automata. *International Journal of Geographical Information Science*, 29, 349-374.
- BROAD, D. R., DANDY, G. C. & MAIER, H. R. 2015. A systematic approach to determining metamodel scope for risk-based optimization and its application to water distribution system design. *Environmental Modelling & Software*, 69, 382-395.
- BROWN, D. G., PAGE, S., RIOLI, R., ZELLNER, M. & RAND, W. 2005. Path dependence and the validation of agent-based spatial models of land use. *International Journal of Geographical Information Science*, 19, 153-174.
- CAO, K., HUANG, B., LI, M. & LI, W. 2014. Calibrating a cellular automata model for understanding rural-urban land conversion: a Pareto front based multi-objective optimization approach. *International Journal of Geographical Information Science*, 28, 1028-1046.
- CAO, M., TANG, G. A., SHEN, Q. & WANG, Y. 2015. A new discovery of transition rules for cellular automata by using cuckoo search algorithm. *International Journal of Geographical Information Science*, 29, 806-824.
- CHAUDHURI, G. & CLARKE, K. C. 2013a. How does land use policy modify urban growth? A case study of the Italo-Slovenian border. *Journal of Land Use Science*, 8, 443-465.
- CHAUDHURI, G. & CLARKE, K. C. 2013b. The SLEUTH land use change model: A review. *The International Journal of Environmental Resources Research*, 1, 88-104.
- CLARKE-LAUER, M. D. & CLARKE, K. C. 2011. Evolving simulation modeling: Calibrating SLEUTH using a genetic algorithm. *Geocomputation*. University College London, London, UK.
- CLARKE, K. C. 2018. Land Use Change Modeling with SLEUTH: Improving Calibration with a Genetic Algorithm. In: CAMACHO OLMEDO, M. T., PAEGELOW, M., MAS, J.-F. & ESCOBAR, F. (eds.) *Geomatic Approaches for Modeling Land Change Scenarios*. Cham: Springer International Publishing.

- CLARKE, K. C., HOPPEN, S. & GAYDOS, L. 1997. A self-modifying cellular automaton model of historical urbanization in the San Francisco Bay area. *Environment and Planning B: Planning and Design*, 24, 247-261.
- CONGALTON, R. G. 1991. A review of assessing the accuracy of classifications of remotely sensed data. *Remote Sensing of Environment*, 37, 35-46.
- CONNOR, J. D., BRYAN, B. A., NOLAN, M., STOCK, F., GAO, L., DUNSTALL, S., GRAHAM, P., ERNST, A., NEWTH, D., GRUNDY, M. & HATFIELD-DODDS, S. 2015. Modelling Australian land use competition and ecosystem services with food price feedbacks at high spatial resolution. *Environmental Modelling & Software*, 69, 141-154.
- CORDER, G. W. & FOREMAN, D. I. 2014. *Nonparametric Statistics*, Somerset, UNITED STATES, John Wiley & Sons, Incorporated.
- COUCLELIS, H. 1985. Cellular worlds: a framework for modeling micro-macro dynamics. *Environment and Planning A*, 17, 585-596.
- COUCLELIS, H. 1989. Macrostructure and microbehaviour in a metropolitan area. *Environment and Planning B: Planning and Design*, 16, 141-154.
- DEB, K., PRATAP, A., AGARWAL, S. & MEYARIVAN, T. 2002. A fast and elitist multiobjective Genetic Algorithm NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6, 182-197.
- ENGELEN, G. & WHITE, R. 2008. Validating and Calibrating Integrated Cellular Automata Based Models of Land Use Change. In: ALBEVERIO, S., ANDREY, D., GIORDANO, P. & VANCHERI, A. (eds.) *The Dynamics of Complex Urban Systems*. Physica-Verlag HD.
- ENGELEN, G., WHITE, R. & DE NIJS, T. C. M. 2003. Environment Explorer: Spatial support system for the integrated assessment of socio-economic and environmental policies in the Netherlands. *Integrated Assessment*, 4, 97-105.
- FENG, Y., LIU, Y., TONG, X., MIAOLONG, L. & DENG, S. 2011. Modelling dynamic urban growth using cellular automata and particle swarm optimization rules. *Landscape and Urban Planning*, 102, 188-196.
- FURTADO, B. A., ETTEMA, D., MACHADO RUIZ, R., HURKENS, J. & VAN DELDEN, H. 2012. A cellular automata intra-urban model with prices and income-differentiated actors *Environment and Planning B: Planning and Design* 39, 897-924.
- GARCÍA, A. M., SANTÉ, I., BULLÓN, M. & CRECENTE, R. 2012. A comparative analysis of cellular automata models for simulation of small urban areas in Galicia, NW Spain. *Computers, Environment and Urban Systems*, 36, 291-301.
- GARCÍA, A. M., SANTÉ, I., BULLÓN, M. & CRECENTE, R. 2013. Calibration of an urban cellular automaton model by using statistical techniques and a genetic algorithm. Application to a small urban settlement of NW Spain. *International Journal of Geographical Information Science*, 27, 1593-1611.
- GIBBS, M. S., MAIER, H. R. & DANDY, G. C. 2012. A generic framework for regression regionalization in ungauged catchments. *Environmental Modelling & Software*, 27-28, 1-14.
- GUO, D., WESTRA, S. & MAIER, H. R. 2017. An inverse approach to perturb historical rainfall data for scenario-neutral climate impact studies. *Journal of Hydrology*, In Press.
- HAASE, D., HAASE, A., KABISCH, N., KABISCH, S. & RINK, D. 2012. Actors and factors in land-use simulation: The challenge of urban shrinkage. *Environmental Modelling & Software*, 35, 92-103.
- HADKA, D. & REED, P. M. 2012. Diagnostic assessment of search controls and failure modes in many-objective evolutionary optimization. *Evolutionary Computation*, 20, 423-452.

- HAGEN-ZANKER, A. 2008. *Measuring the performance of geosimulation models by map comparison*.
- HAGEN-ZANKER, A. 2009. An improved Fuzzy Kappa statistic that accounts for spatial autocorrelation. *International Journal of Geographical Information Science*, 23, 61-73.
- HAGEN-ZANKER, A. & LAJOIE, G. 2008. Neutral models of landscape change as benchmarks in the assessment of model performance. *Landscape and Urban Planning*, 86, 284-296.
- HAGOORT, M., GEERTMAN, S. & OTTENS, H. 2008. Spatial externalities, neighbourhood rules and CA land-use modelling. *The Annals of Regional Science*, 42, 39-56.
- HAINES-YOUNG, R., WEBER, J.-L., PARAMO, F., BRETON, F., PRIETO, O. G. & SOUKUP, T. 2006. Land accounts for Europe 1990-2000. EEA, Copenhagen: European Environment Agency (EEA).
- HEWITT, R., VAN DELDEN, H. & ESCOBAR, F. 2014. Participatory land use modelling, pathways to an integrated approach. *Environmental Modelling and Software*, 52, 149-165.
- JAFARNEZHAD, J., SALMANMAHINY, A. & SAKIEH, Y. 2016. Subjectivity versus Objectivity: Comparative Study between Brute Force Method and Genetic Algorithm for Calibrating the SLEUTH Urban Growth Model. *Journal of Urban Planning and Development*, 142, 05015015.
- KANG, D. & LANSEY, K. 2012. Revisiting Optimal Water-Distribution System Design: Issues and a Heuristic Hierarchical Approach. *Journal of Water Resources Planning and Management*, 138, 208-217.
- KINGSTON, G. B., DANDY, G. C. & MAIER, H. R. 2008. AI Techniques for Hydrological Modeling and Management II: Optimization. In: ROBINSON, L. N. (ed.) *Water Resources Research Progress*. New York: Nova Science Publishers, Inc.
- LAMBIN, E. F., TURNER, B. L., GEIST, H. J., AGBOLA, S. B., ANGELSEN, A., BRUCE, J. W., COOMES, O. T., DIRZO, R., FISCHER, G., FOLKE, C., GEORGE, P. S., HOMEWOOD, K., IMBERNON, J., LEEMANS, R., XIUBIN, L., MORAN, E. F., MORTIMORE, M., RAMAKRISHNAN, P. S., RICHARDS, J. F., SKANES, H., STEFFEN, W., STONE, G. D., SVEDIN, U., VELDKAMP, T. A., VOGEL, C. & XU, J. 2001. The causes of land-use and land-cover change: moving beyond the myths. *Global Environmental Change*, 11, 261-269.
- LAUF, S., HAASE, D., HOSTERT, P., LAKES, T. & KLEINSCHMIT, B. 2012. Uncovering land-use dynamics driven by human decision-making – A combined model approach using cellular automata and system dynamics. *Environmental Modelling & Software*, 27–28, 71-82.
- LI, S., JUHÁSZ-HORVÁTH, L., PEDDE, S., PINTÉR, L., ROUNSEVELL, M. D. A. & HARRISON, P. A. 2017. Integrated modelling of urban spatial development under uncertain climate futures: A case study in Hungary. *Environmental Modelling & Software*, 96, 251-264.
- LI, X., LIN, J., CHEN, Y., LIU, X. & AI, B. 2013. Calibrating cellular automata based on landscape metrics by using genetic algorithms. *International Journal of Geographical Information Science*, 27, 594-613.
- LI, X. & YE, H. A. G.-O. 2004. Data mining of cellular automata's transition rules. *International Journal of Geographical Information Science*, 18, 723-744.
- LIAO, J., TANG, L., SHAO, G., QIU, Q., WANG, C., ZHENG, S. & SU, X. 2014. A neighbor decay cellular automata approach for simulating urban expansion based on particle swarm intelligence. *International Journal of Geographical Information Science*, 28, 720-738.

- LIU, X., SUN, R., YANG, Q., SU, G. & QI, W. 2012. Simulating urban expansion using an improved SLEUTH model. *Journal of Applied Remote Sensing*, 6, 1-19.
- MAAS, A., ENGELEN, G., HAGEN-ZANKER, A., DE NIJS, A. C. M., VAN LOON, J., STRAATMAN, B., WHITE, R., ULJEE, I., VAN DER MEULEN, M. & HURKENS, J. 2005. Calibration and validation of the environment explorer.
- MAIER, H. R., KAPELAN, Z., KASPRZYK, J., KOLLAT, J., MATOTT, L. S., CUNHA, M. C., DANDY, G. C., GIBBS, M. S., KEEDWELL, E., MARCHI, A., OSTFELD, A., SAVIC, D., SOLOMATINE, D. P., VRUGT, J. A., ZECCHIN, A. C., MINSKER, B. S., BARBOUR, E. J., KUCZERA, G., PASHA, F., CASTELLETTI, A., GIULIANI, M. & REED, P. M. 2014. Evolutionary algorithms and other metaheuristics in water resources: Current status, research challenges and future directions. *Environmental Modelling & Software*, 62, 271-299.
- MAIER, H. R., KAPELAN, Z., KASPRZYK, J. & MATOTT, L. S. 2015. Thematic issue on Evolutionary Algorithms in Water Resources. *Environmental Modelling & Software*, 69, 222-225.
- MCGARIGAL, K. 2014. *FRAGSTATS Help*, University of Massachusetts, Amherst.
- MURRAY-RUST, D., BROWN, C., VAN VLIET, J., ALAM, S. J., ROBINSON, D. T., VERBURG, P. H. & ROUNSEVELL, M. 2014a. Combining agent functional types, capitals and services to model land use dynamics. *Environmental Modelling & Software*, 59, 187-201.
- MURRAY-RUST, D., ROBINSON, D. T., GUILLEM, E., KARALI, E. & ROUNSEVELL, M. 2014b. An open framework for agent based modelling of agricultural land use change. *Environmental Modelling & Software*, 61, 19-38.
- MURRAY-RUST, D., ROBINSON, D. T., GUILLEM, E., KARALI, E. & ROUNSEVELL, M. 2014c. An open framework for agent based modelling of agricultural land use change. *Environmental Modelling & Software*, 61, 19-38.
- MUSTAFA, A., HEPPESTALL, A., OMRANI, H., SAADI, I., COOLS, M. & TELLER, J. 2018. Modelling built-up expansion and densification with multinomial logistic regression, cellular automata and genetic algorithm. *Computers, Environment and Urban Systems*, 67, 147-156.
- NEWLAND, C. P., MAIER, H. R., ZECCHIN, A. C., NEWMAN, J. P. & VAN DELDEN, H. 2018a. Multi-objective optimisation framework for calibration of Cellular Automata land-use models. *Environmental Modelling & Software*, 100, 175-200.
- NEWLAND, C. P., ZECCHIN, A. C., MAIER, H. R., VAN DELDEN, H. & NEWMAN, J. P. 2018b. Empirically derived method and software for semi-automatic calibration of Cellular Automata land-use models. *Environmental Modelling & Software*, Under review.
- NGUYEN, D. C. H., DANDY, G. C., MAIER, H. R. & II, J. C. A. 2016. Improved Ant Colony Optimization for Optimal Crop and Irrigation Water Allocation by Incorporating Domain Knowledge. *Journal of Water Resources Planning and Management*, 142, 04016025.
- POGSON, M., RICHARDS, M., DONDINI, M., JONES, E. O., HASTINGS, A. & SMITH, P. 2016. ELUM: A spatial modelling tool to predict soil greenhouse gas changes from land conversion to bioenergy in the UK. *Environmental Modelling & Software*, 84, 458-466.
- PONTIUS JR., R. G., BOERSMA, W., CASTELLA, J.-C., CLARKE, K. C., DE NIJS, T., DIETZEL, C., DUAN, Z., FOTSING, E., GOLDSTEIN, N., KOK, K., KOOMEN, E., LIPPITT, C. D., MCCONNELL, W., MOHD SOOD, A., PIJANOWSKI, B., PITHADIA, S., SWEENEY, S., TRUNG, T. N., VELDKAMP, A. T. & VERBURG,

- P. H. 2008. Comparing the input, output, and validation maps for several models of land change. *The Annals of Regional Science*, 42, 11-37.
- PONTIUS JR., R. G. & PETROVA, S., H. 2010. Assessing the predictive model of land change using uncertain data. *Environmental Modelling and Software*, 25, 299-309.
- REED, P. M., HADKA, D., HERMAN, J. D., KASPRZYK, J. & KOLLAT, J. B. 2013. Evolutionary multiobjective optimization in water resources: The past, present and future. *Advances in Water Resources*, 51, 438-456.
- RIKS 2010. *Map Comparison Kit 3 User Manual*, Maastricht, The Netherlands, The Research Institute of Knowledge Systems.
- RIKS 2011. *MCK Reader: Methods of the Map Comparison Kit*, Maastricht, the Netherlands, The Research Institute of Knowledge Systems.
- RIKS 2015. *Metronamica Documentation*, Maastricht, The Netherlands, the Research Institute of Knowledge Systems.
- RUTLEDGE, D. T., CAMERON, M., ELLIOTT, S., FENTON, T., HUSER, B., MCBRIDE, G., MCDONALD, G., O'CONNOR, M., PHYN, D., POOT, J., PRICE, R., SCRIMGEOUR, F., SMALL, B., TAIT, A., VAN DELDEN, H., WEDDERBURN, M. E. & WOODS, R. A. 2008. Choosing Regional Futures: Challenges and choices in building integrated models to support long-term regional planning in New Zealand\*. *Regional Science Policy & Practice*, 1, 85-108.
- ŞALAP-AYÇA, S., JANKOWSKI, P., CLARKE, K. C., KYRIAKIDIS, P. C. & NARA, A. 2018. A meta-modeling approach for spatio-temporal uncertainty and sensitivity analysis: an application for a cellular automata-based Urban growth and land-use change model. *International Journal of Geographical Information Science*, 32, 637-662.
- SANTÉ, I., GARCÍA, A. M., MIRANDA, D. & CRECENTE, R. 2010. Cellular automata models for the simulation of real-world processes: A review and analysis. *Landscape and Urban Planning*, 96, 108-122.
- SHI, Y.-E., ZUIDGEEST, M., SALZBERG, A., SLIUZAS, R., ZHENG DONG, H., QINGMING, Z., NGUYEN NGOC, Q., HURKENS, J., MINGJUN, P., GUANGHUA, C., MAARSEVEEN, M. V. & DELDEN, H. V. Simulating urban development scenarios for Wuhan. 2012 6th International Association for China Planning Conference (IACP), 17-19 June 2012 2012. 1-13.
- SILVA, E. A. & CLARKE, K. C. 2002. Calibration of the SLEUTH urban growth model for Lisbon and Porto, Portugal. *Computers, Environment and Urban Systems*, 26, 525-552.
- STRAATMAN, B., WHITE, R. & ENGELEN, G. 2004. Towards an automatic calibration procedure for constrained cellular automata. *Computers, Environment and Urban Systems*, 28, 149-170.
- SZEMIS, J. M., MAIER, H. R. & DANDY, G. C. 2012. A framework for using ant colony optimization to schedule environmental flow management alternatives for rivers, wetlands, and floodplains. *Water Resources Research*, 48, n/a-n/a.
- TOBLER, W. R. 1979. Cellular geography. In: GALE, S. & OLSSON, G. (eds.) *Philosophy in Geography*.
- TRUNFIO, G. A. 2006. Exploiting spatio-temporal data for multiobjective optimization of cellular automata models. In: CORCHADO, E., YIN, H., BOTTI, V. & FYFE, C. (eds.) *Intelligent Data Engineering and Automated Learning - IDEAL*. Burgos, Spain.
- VAN DELDEN, H., DÍAZ-PACHECO, J., SHI, Y.-E. & VAN VLIET, J. 2012. Calibration of cellular automata based land use models: lessons learnt from practical experience. *Cellular Automata Modelling for Urban and Spatial Systems*. Oporto, Portugal.

- VAN DELDEN, H. & HURKENS, J. 2011. A generic Integrated Spatial Decision Support System for urban and regional planning. *19th International Congress on Modelling and Simulation - MODSIM*. Perth, Australia.
- VAN DELDEN, H., LUJA, P. & ENGELEN, G. 2007. Integration of multi-scale dynamic spatial models of socio-economic and physical processes for river basin management. *Environmental Modelling and Software*, 22, 223-238.
- VAN DELDEN, H., STUCZYNSKI, T., CIAIAN, P., PARACCHINI, M. L., HURKENS, J., LOPATKA, A., SHI, Y.-E., PRIETO, O. G., CALVO, S., VAN VLIET, J. & VANHOUT, R. 2010. Integrated assessment of agricultural policies with dynamic land use change modelling. *Ecological Modelling*, 221, 2153-2166.
- VAN DELDEN, H., VAN VLIET, J., RUTLEDGE, D. T. & KIRKBY, M. J. 2011. Comparison of scale and scaling issues in integrated land-use models for policy support. *Agriculture, Ecosystems and Environment*, 142, 18-28.
- VAN DELDEN, H. & VANHOUT, R. 2018. A Short Presentation of Metronamica. In: CAMACHO OLMEDO, M. T., PAEGELOW, M., MAS, J.-F. & ESCOBAR, F. (eds.) *Geomatic Approaches for Modeling Land Change Scenarios*. Cham: Springer International Publishing.
- VAN VLIET, J., BREGT, A. K., BROWN, D. G., VAN DELDEN, H., HECKBERT, S. & VERBURG, P. H. 2016. A review of current calibration and validation practises in land-change modeling. *Environmental Modelling & Software*, 82, 174-182.
- VAN VLIET, J., BREGT, A. K. & HAGEN-ZANKER, A. 2011. Revisiting Kappa to account for change in the accuracy assessment of land-use change models. *Ecological Modelling*, 222, 1367-1375.
- VAN VLIET, J., HAGEN-ZANKER, A., HURKENS, J. & VAN DELDEN, H. 2013a. A fuzzy set approach to assess the predictive accuracy of land use simulations. *Ecological Modelling*, 261-262, 32-42.
- VAN VLIET, J., NAUS, N., LAMMEREN, R. J. A., BREGT, A. K., HURKENS, J. & VAN DELDEN, H. 2013b. Measuring the neighbourhood effect to calibrate land use models. *Computers, Environment and Urban Systems*, 41, 55-64.
- VEERBEEK, W., PATHIRANA, A., ASHLEY, R. & ZEVENBERGEN, C. 2015. Enhancing the calibration of an urban growth model using a memetic algorithm. *Computers, Environment and Urban Systems*, 50, 53-65.
- VERBURG, P. H., RITSEMA VAN ECK, J. R., DE NIJS, T. C. M., DIJST, M. J. & SCHOT, P. 2004. Determinants of land use change patterns in the Netherlands. *Environment and Planning B: Planning and Design*, 31, 125-150.
- WANG, F., HASBANI, J.-G., WANG, X. & MARCEAU, D. J. 2011a. Identifying dominant factors for the calibration of a land-use cellular automata model using Rough Set Theory. *Computers, Environment and Urban Systems*, 35, 116-125.
- WANG, F., HASBANI, J.-G., WANG, X. & MARCEAU, D. J. 2011b. Identifying dominant factors for the calibration of a land-use cellular automata model using Rough Set Theory. *Computers, Environment and Urban Systems*, 35, 116-125.
- WANG, Q., GUIDOLIN, M., SAVIC, D. & KAPELAN, Z. 2015. Two-Objective Design of Benchmark Problems of a Water Distribution System via MOEAs: Towards the Best-Known Approximation of the True Pareto Front. *Journal of Water Resources Planning and Management*, 141, 04014060.
- WELCH, B. L. 1947. THE GENERALIZATION OF 'STUDENT'S' PROBLEM WHEN SEVERAL DIFFERENT POPULATION VARIANCES ARE INVOLVED. *Biometrika*, 34, 28-35.
- WHITE, R. & ENGELEN, G. 1993a. Cellular dynamics and GIS: modelling spatial complexity. *Geographical Systems*, 1, 237-253.

- WHITE, R. & ENGELEN, G. 1993b. Complex dynamics and fractal urban form. *In: NIJKAMP, P. & REGGIANI, A. (eds.) Nonlinear Evolution of Spatial Economic Systems*. Berlin: Springer.
- WHITE, R. & ENGELEN, G. 1993c. Fractal urban land-use patterns: a cellular automata approach. *Environment and Planning A*, 25, 1175-1199.
- WHITE, R., ENGELEN, G. & ULJEE, I. 1997. The use of constrained cellular automata for high-resolution modelling of urban land-use dynamics. *Environment and Planning B: Planning and Design*, 24, 323-343.
- WHITE, R., ENGELEN, G. & ULJEE, I. 2015. *Modeling cities and regions as complex systems: from theory to planning applications*, Cambridge, Massachusetts, The MIT Press.
- WICKRAMASURIYA, R. C., BREGT, A. K., VAN DELDEN, H. & HAGEN-ZANKER, A. 2009. The dynamics of shifting cultivation captured in an extended Constrained Cellular Automata land use model. *Ecological Modelling*, 220, 2302-2309.
- WU, F. 2002. Calibration of stochastic cellular automata: the application to rural-urban land conversions. *International Journal of Geographical Information Science*, 16, 795-818.
- WU, W., SIMPSON, A. R., MAIER, H. R. & MARCHI, A. 2012. Incorporation of Variable-Speed Pumping in Multiobjective Genetic Algorithm Optimization of the Design of Water Transmission Systems. *Journal of Water Resources Planning and Management*, 138.
- ZECCHIN, A. C., SIMPSON, A. R., MAIER, H. R., MARCHI, A. & NIXON, J. B. 2012. Improved understanding of the searching behavior of ant colony optimization algorithms applied to the water distribution design problem. *Water Resources Research*, 48, n/a-n/a.
- ZHENG, F., ZECCHIN, A. C., MAIER, H. R. & SIMPSON, A. R. 2016. Comparison of the Searching Behavior of NSGA-II, SAMODE, and Borg MOEAs Applied to Water Distribution System Design Problems. *Journal of Water Resources Planning and Management*, 0, 04016017.
- ZITZLER, E. 1999. *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*. Doctor of Technical Sciences, Swiss Federal Institute of Technology.