Fusion of Correlated Information in Multifidelity Aircraft Design Optimization

by

Zipeng Zhao

B.S., California Institute of Technology (2013)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

Signature redacted

Author ....
Department of Aeronautics and Astronautics
May 21, 2015

Signature redacted

Certified by...........................
Karen Willcox
Professor, Department of Aeronautics and Astronautics
Thesis Supervisor

Signature redacted

Accepted by.........................
Paulo C. Lozano
Chair, Graduate Program Committee

Fusion of Correlated Information in Multifidelity Aircraft Design Optimization

by

Zipeng Zhao

Submitted to the Department of Aeronautics and Astronautics
on May 21, 2015, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

Models used in engineering design often face trade-offs between computational cost and prediction uncertainty. To ameliorate this problem, correlated models of varying fidelities are used together under different fidelity management strategies to produce accurate predictions while avoiding typically expensive costs. However, existing strategies either account for model correlation and operate under the assumption of a strict fidelity hierarchy, or do not consider model correlation but allow model fidelities to vary across the design space. In this thesis, we present a surrogate-based multifidelity framework that simultaneously accounts for model correlation and accommodates non-hierarchical fidelity specifications.

The development of our multifidelity framework can be classified into three stages. The first stage involves the construction of three separate wing weight estimation models that simplify different aspects of the wing sizing problem, thereby creating a scenario where model fidelities are not confined to a rigid hierarchy. The second stage involves the establishment of a formal definition of model correlation, and an extension that allows model correlations to vary across the design space. The third stage involves the incorporation of model correlation in surrogate-based information fusion. To illustrate the application of our framework, we set up a wing weight estimation problem using wing span as design variable. In a later chapter, the problem is extended to two dimensions for increased complexity using body weight and aspect ratio as design variables.

Results from both wing weight estimation problems indicate a combination of variance reduction and inflation at different positions in the design space when model correlation is considered, in comparison to the case where model correlation is ignored.

Thesis Supervisor: Karen Willcox
Title: Professor, Department of Aeronautics and Astronautics

Acknowledgments

My sincere gratitude extends first to Professor Karen Willcox, who offered me the chance to continue my graduate studies through the financial support of a research-assistantship. Professor Willcox also provided continual guidance and encouragement throughout my time of working with her, and helped me significantly refine my research process. I can recall many instances when I faced seeming dead-ends during the project, and she would impart me with insights that led to new perspectives and approaches. There is no doubt about the passion she has for her work.

The research would also not have been possible without the support of Sergio Amaral, a PhD candidate under Professor Willcox. He has always been enthusiastic in answering my questions, and has consistently enlightened me to new takes on problems. Sergio helped me especially with deciphering the mechanisms behind information fusion. He suggested many ways previously unbeknownst to me in exploring matrices. Sergio is also quite entertaining to chat with at the dinner table.

I would also like to thank Rémi Lam, another PhD student of Professor Willcox. He recommended countless resources during my induction to the lab to catch me up to speed. Much of this work is based on Rémi's master's thesis, and he has been very open in clarifying any obscurities I had regarding his multifidelity framework. My gratitude also extends to Chi Feng, a fellow labmate, who helped me understand Gaussian Process Regression. Chi has been a motivating and supportive friend throughout all the six years I have known him. In lab, he never fails to blow our minds with new LaTeX tricks and physics facts. Others that helped me in this project include Ralf Zimmermann, Zheng Wang, Alessio Spantini, Max Opgenoord, and Giulia Pantalone, all of whom work in the Aerospace Computational Design Laboratory.

Lastly, I am immensely grateful for my parents, Li Mu and Yong Zhao. The unbreakable determination and unyielding persistence they showed me as we immigrated across the world gave me the strength I needed to pull through many rough patches along the path of graduate school.

This work was supported by the NASA LEARN program, grant number NNX14AC73A.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Models used in engineering design often face trade-offs between computational cost and prediction uncertainty. In this thesis, a multifidelity framework based on fusion of correlated information is presented. Section 1.1 overviews the existing multifidelity methods, Section 1.2 discusses Bayesian regression techniques for generating surrogates, Section 1.3 explains the motivation behind this endeavor, and Section 1.4 summarizes the research objectives.

1.1 Overview of Multifidelity Methods

Multifidelity design optimization is the concept that information from multiple sources with varying degrees of confidence are combined when searching for the best design as governed by some objective function. In doing so, we harness the inexpensive nature of low order models while achieving the accuracy associated with highly detailed models. With regard to model fidelity, many state-of-the-art frameworks operate under the assumption that some rigid hierarchical structure exists among the models. For example, in the fluid mechanics community, a hierarchy often manifests in the form of Direct Numerical Simulation (DNS), Large Eddy Simulation (LES), Reynolds Averaged Navier-Stokes (RANS), and Euler Equations[26], with DNS being of highest fidelity and Euler Equations being of lowest. By similar logic, for models that numerically solve partial differential equations, hierarchies may be established through differences in fineness of mesh. In general, model fidelity is inversely proportional to model inadequacy, where model inadequacy is defined as the discrepancy between model output and the true value it is assigned to estimate.

Strategies to perform multifidelity design optimization have been studied extensively, and some methods make use of response surface surrogates. Under the assumption of hierarchical model fidelities, these methods leverage on minimizing the discrepancy between the output of the highest fidelity model and those of lower fidelities. The associated computational savings thus come from clever sampling schemes and reduced number of expensive model evaluations. For instance, one might refine surrogates from low fidelity models using evaluations from high fidelity models[5]. Alternatively, one can create surrogates approximating the difference between a high fidelity model and a low fidelity model, thereby improving the confidence of the low fidelity model over a large portion of the design space without sacrificing computational cost[19][17]. All of the aforementioned techniques can be extended to multiobjective optimization. Furthermore, trust region methods have also seen application in the context of multifidelity optimization. Work by Alexandrov et al. demonstrated an approach when gradient information is available[1][2], and March and Willcox developed a method in the derivative-free setting[25].

Given the large number of disciplines considered during aircraft design and the great selection of models within each, an approach that grants more versatility to model fidelity relaxes the assumption of fixed hierarchical structure. Instead, it assigns different models different confidence depending on location in the design space. As such, uncertainties of models with coupled physical mechanisms can be more accurately represented. For instance, a group of wing weight estimation models may take identical inputs, but make simplifications to different aspects of their respective sizing routines. The results they produce would have different accuracies across the design space; the best model in one part of the design space may not be the best in another. This example serves to illustrate the nature of several models we built to guide the development of our multifidelity framework.

1.2 Bayesian Regression Surrogates

When using a surrogate to approximate model output, it is important to assess the associated uncertainty. A rigorous measure of uncertainty is essential in justifying both optimality and robustness of a design. Common approaches to building surrogates that integrate uncertainty information often stem from Bayesian regression. In contrast to the frequentist's technique of devising a solution by minimizing some measure of error relative to a set of training points, the Bayesian technique starts with some hypothesis on the form of the data-generating distribution (the prior) and treats the training points as evidence refining that hypothesis (the likelihood). Since Bayesian inference is performed by mixing and scaling probability distributions, a Bayesian solution (the posterior) withholds a measure of confidence. In application to surrogate optimization, the prior and the likelihood are chosen such that the posterior is analytically tractable. This way, no additional computational resources would be required to reproduce the form of the posterior.

Without reliable knowledge on the class of best fitting parametric models in the context of the underlying physical processes, it is difficult to quantify uncertainty accurately. Under a fixed functional form (e.g. linear, quadratic, etc.), any discrepancy between training points and the posterior mean manifests into global variance inflation. However, this behavior does not match our intuition that uncertainty should be lower where training points are concentrated, and higher where no model evaluation is present to guide the response surface. To address this, Gaussian Process Regression (GPR), also known as Kriging, is adopted as a surrogate construction technique[27][11]. Gaussian Process Regression is a powerful nonlinear regression technique that allows a wide range of functional specifications for the behavior of the response surface and variance in relation to distance from training points. Gaussian Process Regression belongs to the larger class of kernel methods, which has well-known associations in finance, bioinformatics, and artificial intelligence.

Different kernel specifications lead to different surrogate characteristics. In order to enforce smoothness and accommodate gradient-based optimization algorithms, Gaussian Process Regression used in multifidelity frameworks often implement the generalized squared exponential kernel function. If, in addition, gradient information is available, an extension to the Gaussian Process Regression known as Co-Kriging can be applied to build multifidelity surrogates[7][8][14].

Given that Bayesian regression provides only a measure of uncertainty from sampling, we must find a way to account for uncertainty from model inadequacy as well. Our work follows a method proposed by Lam[22], in which the total surrogate variance is broken down into sampling variance and fidelity variance. Sampling variance is the variance computed directly from performing Bayesian regression on the training set, and fidelity variance is a function defined within the design space representing uncertainty in the model's output. Unlike sampling variance, fidelity variance comes from expert elicitation, and is subjective.

1.3 Motivation

Models estimating the same quantity of interest likely operate under similar mechanisms. Consequently, information from similar models are somewhat correlated. Co-Kriging makes use of this correlation by integrating it into covariance matrices[14]. However, Co-Kriging operates under the auto-regressive assumption[20], which necessitates that higher fidelity models encompass all information from lower fidelity models. Therefore, multifidelity frameworks implementing Co-Kriging lack the versatility of non-hierarchical frameworks. Through a separate approach, Lam, Allaire, and Willcox recently developed a multifidelity framework that does not rely on the auto-regressive assumption, but does not explicitly take into account model correlation[23].

When models are correlated, treating their information as independent may result in greater variance reduction than is justified. On the flip side, exploiting correlation may allow us to infer more information from similarities between models than what each model contributes independently. Hence, we propose an extension to the framework developed by Lam, Allaire, and Willcox that takes advantage of model correlation structure in information fusion.

1.4 Objectives

In this thesis, we propose a definition for model correlation that is derived from model discrepancy trends across the design space. We then incorporate model correlation into information fusion and study its effects on the fused surrogate. Finally, we demonstrate the application of the proposed methodology on a multifidelity wing sizing problem.

The thesis is composed of 6 chapters. Chapter 2 provides an overview of the models we created to guide the development of the multifidelity framework. Chapter 3 defines model correlation and applies the idea to the demonstration problem. Chapter 4 conducts information fusion and compares the case in which models are correlated against the case in which models are assumed to be independent. Chapter 5 applies the multifidelity framework to a 2-D wing sizing problem, and Chapter 6 concludes the thesis with a summary and brief discussions of future work.

Chapter 2

Overview of Models

This chapter introduces the data-generating models used to guide the development of the multifidelity framework and presents the setup of the 1-D demonstration problem. Section 2.1 covers the models, and Section 2.2 details the problem formulation.

2.1 Description of Models

This section provides descriptions of four models we used to guide the development of the multifidelity framework.

2.1.1 Transport Aircraft System Optimization (TASOPT)

Transport Aircraft System Optimization (TASOPT) is a physics-based conceptual aircraft design tool developed by Professor Mark Drela at MIT. The software conducts sizing and optimization of a proposed architecture specified through a text-based input file and computes weight, aerodynamic, and engine-performance predictions using fundamental theory and associated computational methods [10]. TASOPT is intended for designing aircraft with potentially unconventional architectures, but is also capable of conducting analysis on existing airframes. Our project makes extensive use of the latter functionality.

TASOPT is a multidisciplinary tool. It captures coupling interactions by iterating over disciplinary subroutines. An example of such process flow can be set up as follows: a subroutine governing aerodynamics computes lift and drag, and feeds the intermediate results to a structural subroutine that computes stresses together with the required geometry at various span-wise positions. The structural subroutine then feeds the calculations into a function that updates the associated weights, perturbing the formerly satisfactory parameters with respect to mission profile and landing load constraints. A propulsion subroutine then revises engine parameters, which, together with current airframe geometry, are returned to the aerodynamics subroutine for updates to lift and drag. This process is repeated until all the constraints are satisfied and the design converges, or until some specified maximum number of iterations has been reached. A more detailed iteration procedure can be found in [9].

This thesis presents an application of our proposed multifidelity framework in a single discipline—wing structures. Nonetheless, the idea can be generalized into multiple coupled disciplines with minimal changes. In order to gain deeper insight into the mechanisms behind TASOPT's structures module, we made slight modifications to the fortran code of the software in order to inspect certain intermediate variables.

In TASOPT, there are two components to wing weight; they are wing box weight and add-ons weight respectively. Wing box weight is further broken down into inboard weight and outboard weight, with inboard representing the wing segment between the fuselage and the engine, and outboard representing the segment extending beyond the engine. Add-ons weight encompasses skins, control surfaces, and other minor structural elements, and is computed based on an empirical ratio relative to total wing weight.

Wing box is sized according to moment and shear loads at wing root and wing break. Wing root defines the position where the inboard intersects the fuselage, while wing break defines the position where the engine is attached. Wing box consists of three key geometric parameters - spar cap thickness (t_{cap}), spar web thickness (t_{web}), and chord length (c). Spar cap is sized according to wing root and break bending moments, spar web is sized according to wing root and break shear loads, and chord length is assumed to follow a triangular distribution along wing span, constrained by taper ratios. Wing box height (h_{box}) and width (w_{box}) are inferred based on ratios relative to local chord, and are adjustable in the input file. Figure A-5 is taken from TASOPT documentation [10], and offers a visual reference of the wing box. TASOPT does not allow wing span to be manually specified, and instead sizes for it based on required lift in limit operating cases. As a result, higher cruise altitude typically correspond to higher wing spans. We will take advantage of this relationship later in setting up the multifidelity problem. Additional details on TASOPT's wing sizing routines can be found in [10].

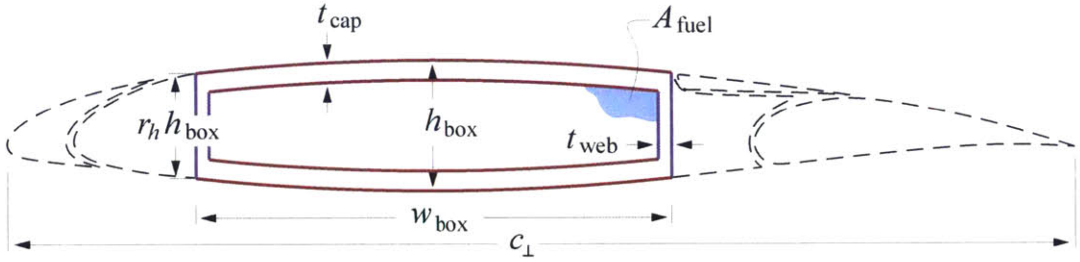

Figure 2-1: The solid lines represent the wing box as modeled by TASOPT, while the dotted lines outline the general shape of the wing. The figure is taken from [10].

2.1.2 Finite Element Model (FEM)

We develop a one-dimensional finite element model (FEM) as a simple model of the wing structure. The model takes in geometric parameters directly and does not size the wing internally. The FEM assumes a hollow rectangular beam as an approximation to the shape of the wing box, and computes secondary component weights using a similar technique as TASOPT. The FEM's functionality includes both wing weight prediction and wing tip deflection estimation, although we are only interested in using the former for this exercise.

We consider the FEM to have model inadequacy in wing weight estimation. The rough approximation to wing geometry, along with the lack of an iterative sizing routine likely makes the tool inaccurate for computing wing weight. Table 2.1 summarizes the input to the model. Note that in the table, as well as throughout the thesis, we use "weight" and "mass" interchangeably. Both terms carry units of kg.

Variable Name	Units	Description		Variable Name	Units	Description
E	$[N/m^2]$	Young's Modulus		t_{cap}	$[m]$	Spar Cap Thickness
ρ	$[kg/m^3]$	Material Density		t_{web}	$[m]$	Spar Web Thickness
b	$[m]$	Wing Span		m_{body}	$[kg]$	Empty + Payload - Wing - Engines Mass
h_{wbox}	$[m]$	Wing Box Height		m_{eng}	$[kg]$	Engine Mass
w_{box}	$[m]$	Wing box Width		N	$[\cdot]$	Critical Loading Factor

Table 2.1: The columns summarize input parameters to the FEM.

To compute wing weight, the FEM multiplies the volume of the wing box with material density, then adds secondary component weights by assuming that they account for 31.5% of total wing weight. The volume of the wing box is given by span-weighted averages of t_{cap}, t_{web}, h_{wbox}, w_{box}, and c from TASOPT.

2.1.3 Rectangular Wing Weight Model (RWW)

The Rectangular Wing Weight Model (RWW) is another simple model used to estimate wing weight. It contains an internal sizing routine that implements similar physical principles as TASOPT. The RWW assumes a rectangular wing profile, and computes secondary component weights in the same way as TASOPT. Table 2.2 summarizes the input to the RWW.

Variable Name	Units	Description		Variable Name	Units	Description
AR	$[\cdot]$	Aspect Ratio		m_{engine}	$[m]$	Engine Mass
b	$[m]$	Wing Span		t/c	$[m]$	Thickness/Chord Ratio
m_{body}	$[kg]$	Empty + Payload - Wing - Engines Mass		N	$[\cdot]$	Critical Loading Factor
m_{fuel}	$[kg]$	Fuel Mass				

Table 2.2: The columns summarize input parameters to the RWW.

The model first guesses a wing weight, then calculates the moment and shear loads at wing root using equations 2.1 and 2.2,

$$
\begin{aligned}
S_o =& \frac{m_{body}\, gN}{2} && (2.1)\\
M_o =& \frac{m_{body}\, gN}{2} \cdot \frac{b}{2} \cdot \eta_{lift} + \frac{m_{eng} gN}{1} \cdot \frac{b}{2} \cdot \eta_{lift} && (2.2)\\
& - \frac{m_{eng}\, gN}{1} \cdot \frac{b}{2} \cdot \eta_s \\
& - \frac{m_{wings}\, gN}{2} \cdot \frac{b}{2} \cdot (\eta_{cm} - \eta_{lift}) \\
& - \frac{m_{fuel}\, gN}{2} \cdot \frac{b}{2} \cdot (\eta_{cm} - \eta_{lift}),
\end{aligned}
$$

where S_o and M_o represent shear and moment at wing root respectively, η_{lift} represents center-of-lift as a fraction along semi-span (with 0 taken to be centerline of fuselage and 1 taken to be wing tip), and η_{cm} represents wing and fuel tank center-of-mass as a fraction along semi-span. In particular, the RWW assumes a near-elliptic spanwise lift distribution along the wing such that $\eta_{lift} = \frac{4}{3\pi} + 0.01$ (outward shift of 0.01 from true elliptic lift distribution). This is a reasonable approximation given the assumed geometry of the wing [4]. The assumption of rectangular wing box translates to a wing and fuel tank center-of-mass at $\eta_{cm} = 0.5$, or exactly halfway across the semi-span.

After obtaining the shear and moment loads, the RWW uses manually specified maximum acceptable bending and shear stresses to size the spar cap and spar web respectively. These numbers are chosen to be consistent with those used in TASOPT. Equations 2.3 and 2.4 compute spar cap thickness-over-chord ($\bar{t}_{cap}$) and spar web area (A_{web}) respectively.

$$
\begin{aligned}
\sigma_{cap} &= \frac{M_o}{\bar{c}^3} \cdot \frac{6\bar{h}}{\bar{w}} \cdot \frac{1}{\bar{h} - (\bar{h} - 2\bar{t}_{cap})^3} && (2.3)\\
\tau_{web} &= \frac{S_o}{A_{web}} && (2.4)
\end{aligned}
$$

where σ_{cap} and τ_{web} represent the aforementioned maximum acceptable bending and shear stresses, $\bar{c}$ represents standard mean chord length $[m]$, and $\bar{h}$ and $\bar{w}$ represent average height and width of wing box divided by $\bar{c}$.

Once the new wing box geometry parameters are determined, the RWW recalculates m_{wings} and stores its value for possible next iteration. If the current wing weight and the wing weight from previous iteration differs by less than 10 kg, the procedure terminates and returns the current wing weight as output. Otherwise, the RWW proceeds onto the next iteration with the latest wing weight as the new guess. The number of iterations until convergence varies depending on the problem setup; the data in this thesis are produced typically in under 10 iterations.

2.1.4 Elliptic Wing Weight Model (EWW)

The Elliptic Wing Weight Model (EWW) builds upon the RWW and incorporates additional details with regard to lift distribution and wing geometry. As the name implies, the EWW assumes an elliptic wing, which translates to variable chord lengths across the wing span. As in the RWW, the wing spar caps and spar webs are sized based on moment and shear loads at wing root. But unlike the RWW, the root position is taken to be the intersection between the wing and the fuselage rather than at the centerline of the fuselage. An extension to this change manifests in the form of an altered lift distribution, where the segment of the wing inside the fuselage does not provide lift, but the fuselage carries over some of the lift from the exposed portion of the wing. The benefit to computing the exact magnitude of the fuselage carryover lift is quite small given the uncertainty due to other sources of error in this model, so an estimation is made that relies on an adjustable ratio representing the fraction of lift relative to a true wing the fuselage provides. By default, this ratio is set to be consistent with TASOPT ($f_{carryover} = 0.5$). Note that in the latest version of TASOPT, fuselage carryover is accounted for in a different manner, so this thesis references an older version of the software (v. 2.08).

Table 2.3 summarizes the input to the EWW. They are mostly the same as those in the RWW, with the addition of in-fuselage wing segment length (b_o) as a variable.

Due to the elliptic wing geometry and the presence of fuselage carryover, moment at wing root is computed in a different manner from the RWW. Equation 2.5 captures

Variable Name	Units	Description		Variable Name	Units	Description
AR	$[\cdot]$	Aspect Ratio		m_{fuel}	$[kg]$	Fuel Mass
b	$[m]$	Wing Span		m_{eng}	$[kg]$	Engine Mass
b_o	$[m]$	In-Fuselage Wing Segment Length		t/c	$[m]$	Thickness/Chord Ratio
m_{body}	$[kg]$	Empty + Payload - Wing - Engines Mass		N	$[\cdot]$	Critical Loading Factor

Table 2.3: The columns summarize input parameters to the EWW.

the new approach.

$$\begin{aligned} M_o =&(\frac{b}{2})^2(y_{o,lift} - y_{o,weight})[\frac{1}{4}\eta_o\ (2sin^{-1}(\eta_o) + sin(2sin^{-1}(\eta_o)) - \pi) \\ &+ \frac{cos^3(sin^{-1}(\eta_o))}{3}] - (\frac{m_{eng}\ gN}{1} \cdot \frac{b}{2}(\eta_s - \eta_o)) \end{aligned} \tag{2.5}$$

where

$$\begin{aligned} y_{o,lift} =&\frac{4L_o}{b/2}[(-2sin^{-1}(\eta_o) - 2\eta_o\ cos(sin^{-1}(\eta_o)) + \pi) \\ &+ 2f_{carryover}(sin^{-1}(\eta_o) + \eta_o\ cos(sin^{-1}(\eta_o)))]^{-1} \\ y_{o,weight} =&\frac{4w_o}{\pi\ b/2} \\ L_o =&\frac{m_{body}\ gN}{2} + \frac{m_{eng}\ gN}{1} + \frac{m_{wings}\ gN}{2} + \frac{m_{fuel}\ gN}{2} \\ w_o =&\frac{m_{wings}\ gN}{2} + \frac{m_{fuel}\ gN}{2} \end{aligned}$$

The equations for computing moment are derived by integrating the elliptic lift distribution associated with the EWW's wing geometry. Aside from the revised moment, the EWW also numerically integrates for wing weight to accommodate variable chord length along wing span. However, it does not apply appropriate reductions to structural requirements for the wing segment inside the fuselage despite it carrying a mitigated load. This leads to potential overestimate of wing weight if the moment at wing root is high. Thus, the EWW has higher fidelity in certain regions of the design space than the RWW and the FEM, but lower fidelity in others, giving rise to a scenario where our models are not aligned in a strict fidelity hierarchy.

2.1.5 Model Validation

We are interested in modeling the wing weight of medium-range commercial transport aircraft, so we set up our model validation process based on the well-known Boeing 737-800 design. The 737-800 is a variant of the third generation derivative of the 737, and has been in service around the world since 1998. The aircraft features two wing-mounted turbofan engines, a back-swept cantilever wing, and a carrying capacity of up to 189 passengers[13]. In addition, the aircraft is designed for missions with cruise range around 5500 km, cruise altitude around 11000 m, and cruise speed of Mach 0.785[18].

We use TASOPT's estimate of 737-800 performance parameters as a baseline against which to validate. Previous work with TASOPT using the same reference aircraft has been conducted with extensive validation against parameters published through Boeing's website, therefore we are confident in the model's estimates. Table 2.4 presents the result from all our models using shared inputs corresponding to the parameters of the reference aircraft.

The FEM, RWW, and EWW all model the 737-800 reasonably well. In this particular scenario, the FEM performs the best because it obtains wing geometry parameters directly from TASOPT and does not size anything on its own. The FEM assumes a constant wing box profile and uses span-weighted average of c, t_{cap}, t_{web}, etc. to compute wing weight. Since the wing box structural components do not vary linearly across wing span, the FEM often underestimates the inboard weight and overestimates the outboard weight, leading to an accurate overall wing weight only in special cases where the aforementioned errors cancel out. As such, the model works well when we can ensure that the pattern of variation in structural components can be suitably approximated by the weighted averaging method. This limits its predictive power to the class of cantilever wings that are roughly rectangular in geometry, have nearly uniform chord lengths, and carry similar loads at wing root and wing break.

The RWW outperforms the EWW in this setup. A rectangular wing is a better approximation than an elliptic wing for 737-800, and the lift profile for the actual

		Physics-Based Methods				
Variable	**Units**	**TASOPT**	**The FEM**	**The RWW**	**The EWW**	**Torenbeek**
AR	$[\cdot]$	9.15		9.15	9.15	
b	$[m]$	37.816	37.816	37.816	37.816	37.816
c	$[m]$	1.668-6.176		4.133	4.133	
Λ (Sweep)	$[deg]$	25.33		0	0	25.33
S	$[m^2]$	127.6		156.3	156.3	
λ (Taper Ratio)	$[\cdot]$	0.27				
N	$[\cdot]$	3	3	3	3	3
t/c	$[\cdot]$	0.109-0.190		0.156	0.156	
h_{box}	$[m]$	0.182-1.173	0.505	0.580	0.580	
w_{box}	$[m]$	0.807-2.988	1.868	2.066	2.066	
t_{cap}	$[m]$	0.00956-0.02480	0.0165	0.0176	0.0182	
t_{web}	$[m]$	0.00245-0.00636	0.0043	0.0032	0.0032	
M_o	$[Nm]$	4595769		4097204 (-10.8%)	4210875 (-8.4%)	7883000*
m_{eng}	$[kg]$	5708	5708	5708	5708	
m_{fuel}	$[kg]$	22127		22126	22126	
m_{empty}	$[kg]$	56164				56164
$GTOW$	$[kg]$	78291				78291
m_{body}	$[kg]$	35029	35029	35029	35029	
m_{wing}	$[kg]$	9719.0	9835.7 (+1.2%)	11372 (+15.6%)	11800 (+20.0%)	
Input		*** No relief due to engine**				
Output						

Table 2.4: Model validation is performed using 737-800 as reference aircraft. Cells colored in light blue represent input to the particular model, cells colored in orange represent output from the particular model, cells colored in gray represent reference parameters, and cells with no color represent intermediate variables. Note that all models implement physics-based methods for wing weight estimation.

aircraft is closer to that in the RWW than in the EWW. Additionally, the RWW underestimates the moment at wing root to a greater extent than the EWW, which allows it to offset the overestimate in wing spar cap thickness more effectively than the EWW. The overestimate in wing spar cap is a result of using the standard mean chord rather than variable chord lengths in sizing—a characteristic common to both the RWW and the EWW.

The EWW computes the bending moment at every point along the wing. In order to ensure that the model correctly accounts for engine relief and that the moment distribution matches our intuition, we refer to Figures 2-2 and 2-3. Figure 2-2 is given in TASOPT documentation [10], while Figure 2-3 is generated by the EWW.

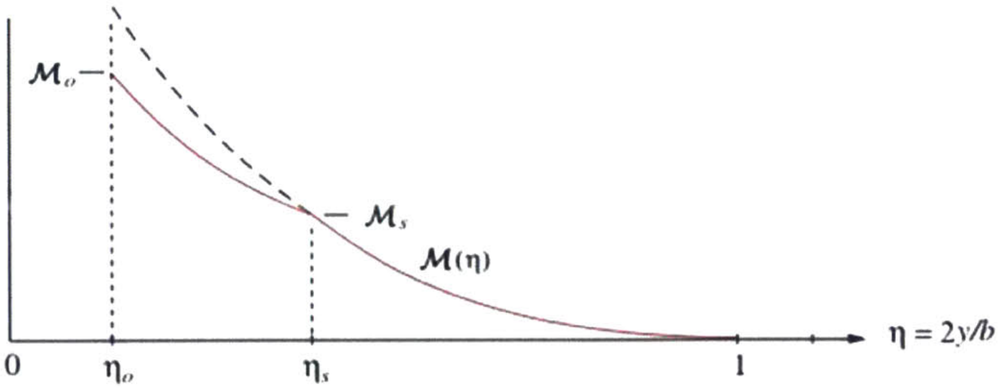

Figure 2-2: The red curve shows the general trend of the spanwise moment distribution according to TASOPT. η_s represents the engine attachment location. The engine weight contributes a substantial amount of moment relief at η_o, where the wing intersects the fuselage. The figure is taken from [10].

Figure 2-2 provides a sketch illustrating how moment at wing root should be modified due to engine attachment. We observe similar behavior in the EWW. From Table 2.4, we notice that Torenbeek's wing root bending moment exceeds those from the other models. The formula used by Torenbeek is presented in Equation 2.6[29]

$$M_o = \frac{1}{4} b \; N \; \eta_{lift} \; GTOW[1 - \frac{\eta_{fuel}}{\eta_{lift}}(1 - \frac{m_{empty}}{GTOW})], \tag{2.6}$$

where η_{fuel} and η_{lift} represent coordinate of center of mass for fuel tank and center of lift respectively, both relative to semi-span. Specifically, we assume the fuel tank to coincide with the wing box ($\eta_{fuel} = \eta_{cm}$), and the lift to follow an elliptic profile ($\eta_{lift} = \frac{4}{3\pi}$). Torenbeek's overestimation stems from the absence of moment relief due to engine attachment. The EWW accounts for that moment relief, but fails to trim down the structural requirement for the in-fuselage segment of the wing. As a result, the EWW properly computes moment at wing root, but overestimates the overall wing weight.

The models we use all compute wing weight differently, with some incorporating more details than others, and some working under more suitable assumptions than others. It is easy to see that the models are somewhat correlated, and we proceed to show how to take advantage of that by first setting up a demonstration problem.

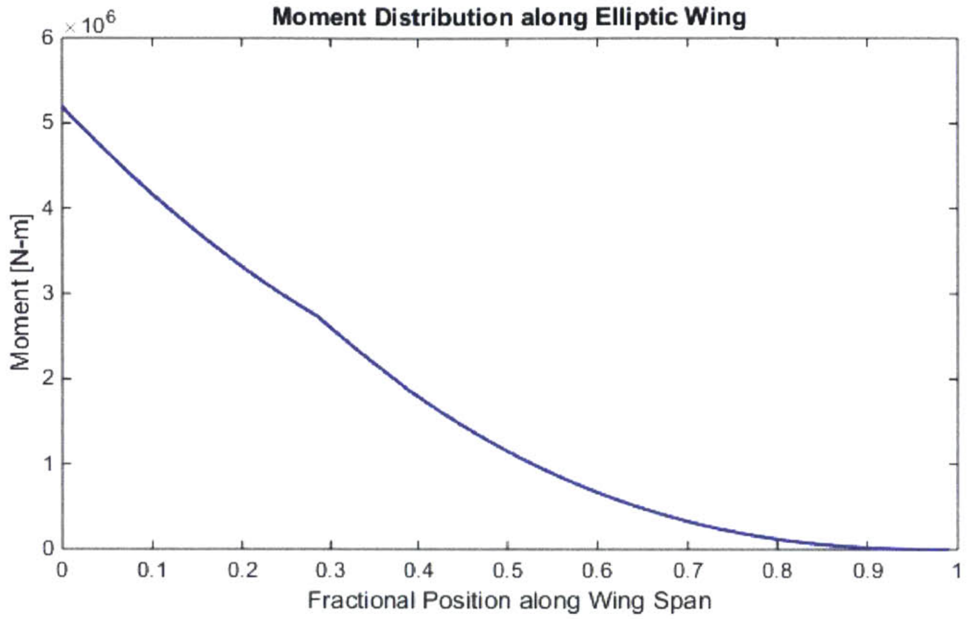

Figure 2-3: The blue curve shows the spanwise moment distribution computed for the 737-800 using the EWW. Rather than terminating at the wing root, the calculation is extended to the portion of the wing inside the fuselage. In this particular setup, η_s is fixed at 0.285, and η_o is taken to be 0.106.

2.2 1-D Problem Setup

This section details the setup of the demonstration problem. The problem serves as a context under which our multifidelity framework is developed.

2.2.1 Reference Aircraft

As in model validation, we adopt the 737-800 as reference aircraft for the demonstration problem. This implies that any parameters not treated as design variables will assume values matching the technical specifications of the reference aircraft. Of course, this is a gross simplification of the true multidisciplinary nature of the design problem. In doing so, we introduce many designs that may not be realistic or even feasible, given the lack of consideration for coupling effects our models should have against other disciplines. Nevertheless, these uncertainties are reflected in model fidelity definitions, which will be discussed in Chapter 3. A more comprehensive list of technical specifications of the reference aircraft can be found in Appendix A. Table 2.5

summarizes the default values of a few parameters most relevant to wing sizing.

Variable Name	Units	Value	Description
m_{body}	$[kg]$	35029	Empty + Payload - Wing - Engines Mass
m_{eng}	$[kg]$	5708	Sized for Cruise at Mach 0.78, 35000 ft
m_{fuel}	$[kg]$	22127	Sized for 5556 km Mission, 180 Passengers
b_o	$[m]$	3.607	Sized for Fuselage Diameter of 3.91 m
b	$[m]$	34.18	Sized for Cruise at Mach 0.78, 35000 ft
c_o	$[m]$	6.176	Chord at Wing Root
c_s	$[m]$	4.323	Chord at Wing Break (Engine Attachment)
c_t	$[m]$	1.668	Chord at Wing Tip
AR	$[\cdot]$	9.15	Aspect Ratio
Λ	$[deg]$	25.33	Fixed Sweep, not Considered by FEM, RWW and EWW
η_s	$[\cdot]$	0.285	Engine Attachment Position along Semi-Span

Table 2.5: The columns provide a subset of technical specifications of the reference aircraft (737-800). A more comprehensive list of technical specifications can be found in Appendix A. All units are given in SI, and brief descriptions are provided alongside each variable.

2.2.2 Design Variable Selection

We would like to explore the relationship between wing weight and wing span over a range of aspect ratios, so we replace the reference b and AR values with those we manually specify. However, TASOPT does not allow b to be chosen explicitly. In order to generate data with TASOPT, we rely on cruise altitude as a tuning knob. Figure 2-4 shows the relationship between the two at fixed aspect ratio. Clearly, we can take advantage of the positive monotonicity between wing span and altitude. A drawback to this method is that we cannot enforce ordered structure to the input vector (e.g. evenly-spaced grid).

Given that aspect ratio is defined by the square of wing span divided by wing area, we expect b and AR to be highly correlated. Figure 2-5 shows the relationship under constant altitude. Typically, it is advisable to select design variables that

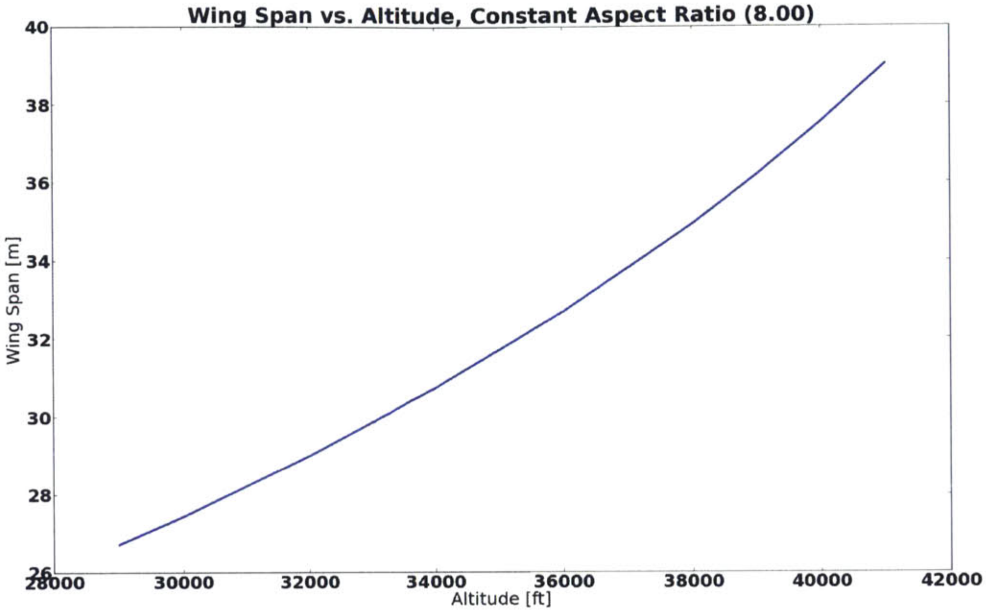

Figure 2-4: The blue curve represents the wing span values converged upon by TASOPT's internal sizing routine as cruise altitude is varied. We take advantage of the monotonicity between the two variables by explicitly specifying altitude to control wing span.

are independent. In this case, only wing span is treated as design variable, while aspect ratio is used to expand the set of data points available for computing model correlation. Therefore, collinearity between wing span and aspect ratio does not introduce unnecessary complications to the setup.

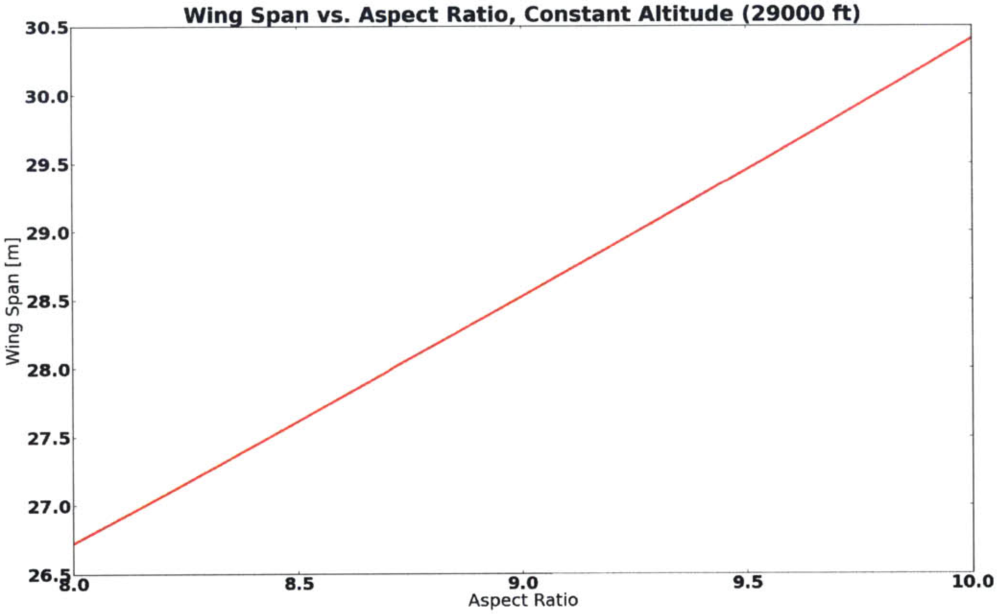

Figure 2-5: The red curve shows the relationship between wing span and aspect ratio for the reference aircraft (737-700) under constant altitude (29000 ft). The linear nature of the curve suggests high collinearity between wing span and aspect ratio. Note that for the demonstration problem, only wing span is used as design variable, while aspect ratio is perturbed to generate more data for characterizing model correlation.

Chapter 3

Model Correlation and Fidelity

This chapter presents a method for quantifying model correlation and model fidelity. Section 3.1 details the process through which we characterize model correlations, and Section 3.2 assigns fidelity variance to respective models using expert elicitation.

3.1 Characterizing Correlation

This section covers the method for computing model correlation, and applies the idea to the demonstration problem.

3.1.1 Defining Model Correlation

Model discrepancy, used interchangeably with model error in this thesis, is defined as the difference between a model's output and the true value it is assigned to estimate. The idea originated with Goldstein[15] and Kennedy[21] in the context of general computer simulations. More formally, model discrepancy can be represented in the following manner. Let $\boldsymbol{d^1}$ be the vector of design variables characterizing design 1, $\boldsymbol{d^2}$ be the vector of design variables characterizing design 2, ..., and $\boldsymbol{d^n}$ be the vector of design variables characterizing design n. Assume without loss of generality that each $\boldsymbol{d^k}$ has length m. Let $z_A|\boldsymbol{d^1}$ be the output of model A evaluated at $\boldsymbol{d^1}$, $z_B|\boldsymbol{d^1}$ be the output of model B evaluated at $\boldsymbol{d^1}$, ..., and $z_{truth}|\boldsymbol{d^1}$ the output of the "truth model" evaluated at $\boldsymbol{d^1}$. All $z_A|\boldsymbol{d^k}$, $z_B|\boldsymbol{d^k}$, and $z_{truth}|\boldsymbol{d^k}$ are scalars. We have,

$$\begin{aligned} D_A^1 &= z_A|\boldsymbol{d^1} - z_{truth}|\boldsymbol{d^1} & D_B^1 &= z_B|\boldsymbol{d^1} - z_{truth}|\boldsymbol{d^1} \\ D_A^2 &= z_A|\boldsymbol{d^2} - z_{truth}|\boldsymbol{d^2} & D_B^2 &= z_B|\boldsymbol{d^2} - z_{truth}|\boldsymbol{d^2} \\ &\vdots & &\vdots \\ D_A^n &= z_A|\boldsymbol{d^n} - z_{truth}|\boldsymbol{d^n} & D_B^n &= z_B|\boldsymbol{d^n} - z_{truth}|\boldsymbol{d^n}, \end{aligned} \tag{3.1}$$

where D_A^k, D_B^k are model discrepancies of A and B based on design k. D_A^k and D_B^k can be treated as random variables. Note that the random nature of the variables comes from the lack of information on $z_{truth}|\boldsymbol{d^k}$. Elements obtained by iterating over k in the sets $\{D_A^k\}$ and $\{D_B^k\}$ form the vectors $\boldsymbol{D_A}$ and $\boldsymbol{D_B}$, both of which have length n.

When we refer to model correlation, we refer explicitly to the degree of linearity in model discrepancy trends. Equation 3.2 gives the precise mathematical definition

$$\rho_{A,B} = \frac{cov(\boldsymbol{D_A}, \boldsymbol{D_B})}{\sqrt{var(\boldsymbol{D_A})var(\boldsymbol{D_B})}} \tag{3.2}$$

where $\rho_{A,B}$ is the correlation between models A and B. Notice that this is exactly the Pearson's correlation coefficient between $\boldsymbol{D_A}$ and $\boldsymbol{D_B}$[28]. When more than two models are involved, the correlation between each pair of models is computed and organized in a symmetric, positive-definite correlation matrix of the form

$$\boldsymbol{\rho} = \begin{bmatrix} 1 & \rho_{\mathcal{M}_1,\mathcal{M}_2} & \cdots & \rho_{\mathcal{M}_1,\mathcal{M}_N} \\ \rho_{\mathcal{M}_2,\mathcal{M}_1} & 1 & \cdots & \rho_{\mathcal{M}_2,\mathcal{M}_N} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{\mathcal{M}_N,\mathcal{M}_1} & \rho_{\mathcal{M}_N,\mathcal{M}_2} & \cdots & 1 \end{bmatrix}$$

where N is the number of models (not including the "truth model") and $\mathcal{M}_i$ enumerates the models.

In application to the demonstration problem, z_A, z_B, and z_C are produced by the FEM, the RWW, and the EWW, respectively. As in model validation, we use TASOPT to approximate true aircraft parameters, so z_{truth} is produced by TASOPT.

Based on past validation studies, we assume TASOPT's wing weight estimates to fall within ±5% relative to true unknown values over all $\boldsymbol{d^k}$. If wing weight data and detailed technical specifications for a large number of aircraft are available, we can use that information directly instead of relying on estimated parameters from TASOPT. However, for the sake of convenience, we adopt TASOPT as the "truth model", against which the FEM, the RWW, and the EWW compute their respective model discrepancies.

Models that work in similar ways tend to exhibit similar discrepancy trends. But similar discrepancy trends do not necessarily imply the same mechanisms behind the models. Hence, it is important that we understand the mechanisms behind the models and are able to explain discrepancy trends through model physics. As such, model correlation is no longer a mere artifact of mathematics, but carries physical significance that can be quite valuable to the design process.

3.1.2 Data for Model Correlation

Different training data can be used to obtain different model discrepancy vectors. Selection of the training data should be done under the engineer's discretion. For instance, if the wing we are interested in has relatively high aspect ratio, we should refrain from including wings with low aspect ratios in the training data. This way, the model correlations will apply more specifically to the type of aircraft we are interested in, despite the models' ability to handle other wing architectures. In other words, instead of weighing all training points evenly, we can impose certain restrictions on the available data so that points most similar to our intended design are weighted higher than those farther away.

In our problem, the training set is generated based on perturbations of two parameters from the reference aircraft, AR and b. AR is sampled by taking 50 evenly-spaced values from 8.0 to 10.0, inclusive, and b is sampled by taking 50 evenly-spaced values of altitude from 29000 ft to 41000 ft, inclusive. In order to generate the design variable vectors, we submitted all combinations of 50 aspect ratios and 50 altitudes to TASOPT and obtained the associated input parameters for the FEM, RWW, and

EWW. Each aforementioned combination corresponds to a unique aircraft and wing design, leading to a total of 2500 unique designs. Furthermore, we allow m_{body} and m_{eng} to vary freely, giving us a training set that covers a wider range of wing architectures. Figure 3-1 shows the discrepancies of the FEM, the RWW, and the EWW projected onto b.

Figure 3-1: Colored points represent model discrepancies based on a set of inputs obtained by perturbing the reference design (737-800). The blue points represent the model discrepancies of the FEM, the green points represent the model discrepancies of the RWW, and the purple points represent the model discrepancies of the EWW.

Based on this training set, the FEM underestimates wing weight, the EWW overestimates wing weight at high wing spans, and the RWW performs consistently everywhere. Differences in these trends owe their nature to different flaws in model physics. In particular, the FEM's method for interpolating wing box geometry through fixed points along semi-span leads to an underestimation of inboard weight that is not cancelled by an overestimation of outboard weight with similar magnitude. Its model discrepancy grows with wing span because the failure to capture non-linearities in wing profile becomes more and more exacerbated as wing span increases. the RWW's assumption of rectangular wing geometry, coupled with an iterative sizing routine that somewhat corrects for segment weight miscalculations makes the model relatively proficient throughout our design space. The EWW's incorporation of fuselage carryover results in more accurate wing weight estimations compared to the RWW at lower spans, but the model overcompensates by failing to reduce the structural requirements for the wing segment inside the fuselage at higher spans. More specifically, at lower spans, the EWW slightly underestimates wing root moment despite fuselage carryover corrections. This underestimation cancels out some of the errors from oversizing the wing segment inside the fuselage, giving the EWW an edge over the RWW. However, at higher spans, the EWW more accurately computes wing root moment, making the oversizing mitigation disappear. The result is the pattern of model discrepancies in Figure 3-1, which should aid us in characterizing model correlations and assigning model fidelities.

3.1.3 Region Selection

Just as models may perform better in one part of the design space than in another, their correlations may also change depending on the location in the design space. To capture these changes, we propose a guideline for stratifying the training data. We begin with a naive method where we divide the design space into equally-sized regions, each containing a unique subset of the training data (i.e. the first region contains $\boldsymbol{d}^1$, $\boldsymbol{d}^2$, ..., $\boldsymbol{d}^p$, the second region contains $\boldsymbol{d}^{p+1}$, $\boldsymbol{d}^{p+2}$, ..., $\boldsymbol{d}^{p+q}$, and so on.). Model correlation associated with each region is characterized using exclusively the training data contained within that region (i.e. the first region uses p points to characterize model correlation, the second region uses q points to characterize model correlation, and so on.). Provided that the total number of data points stays fixed, the more regions there are, the fewer the number of data points per region. The top plots of Figures 3-2, 3-3, 3-4, and 3-5 show the model correlations computed using this approach for different stratification schemes, and the bottom plots show the corresponding distribution of data points across the regions.

Data points tend to become sparse near the edges of the design space, and there can be as few as 2 or 3 if too many regions are chosen (e.g. 75 regions). Regions

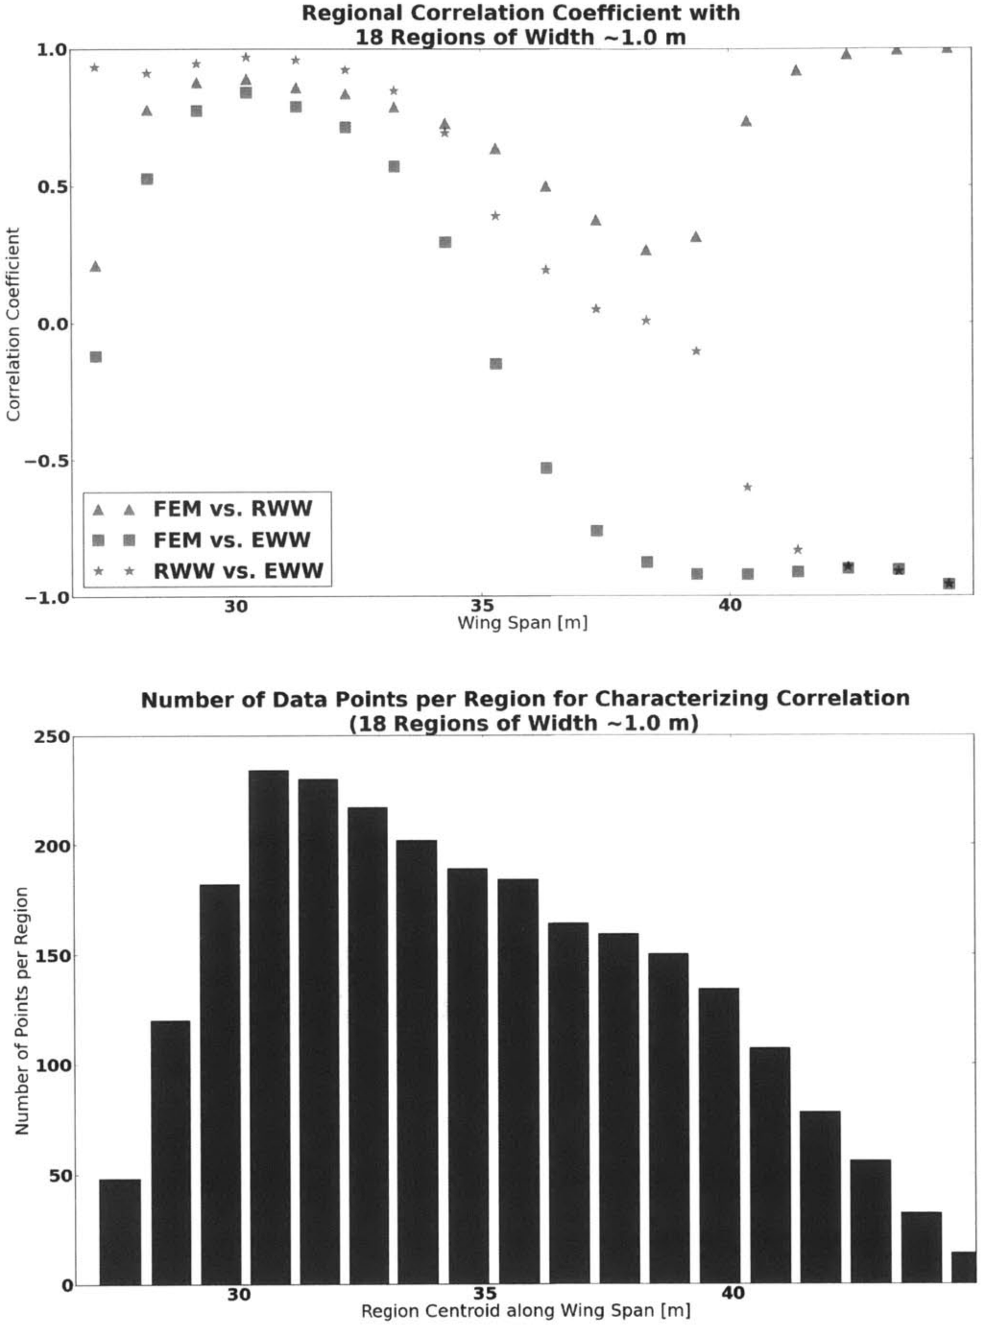

Figure 3-2: The design space is divided into 18 equally-sized regions. Top: the markers represent correlations between different pairs of models at region centroids. Triangles correspond to the correlation between the FEM and the RWW. Squares correspond to the correlation between the FEM and the EWW. Stars correspond to the correlation between the RWW and the EWW. Bottom: the bars correspond the the number of data points used for characterizing correlation. The left edge of each bar coincides with the centroid of its associated region.

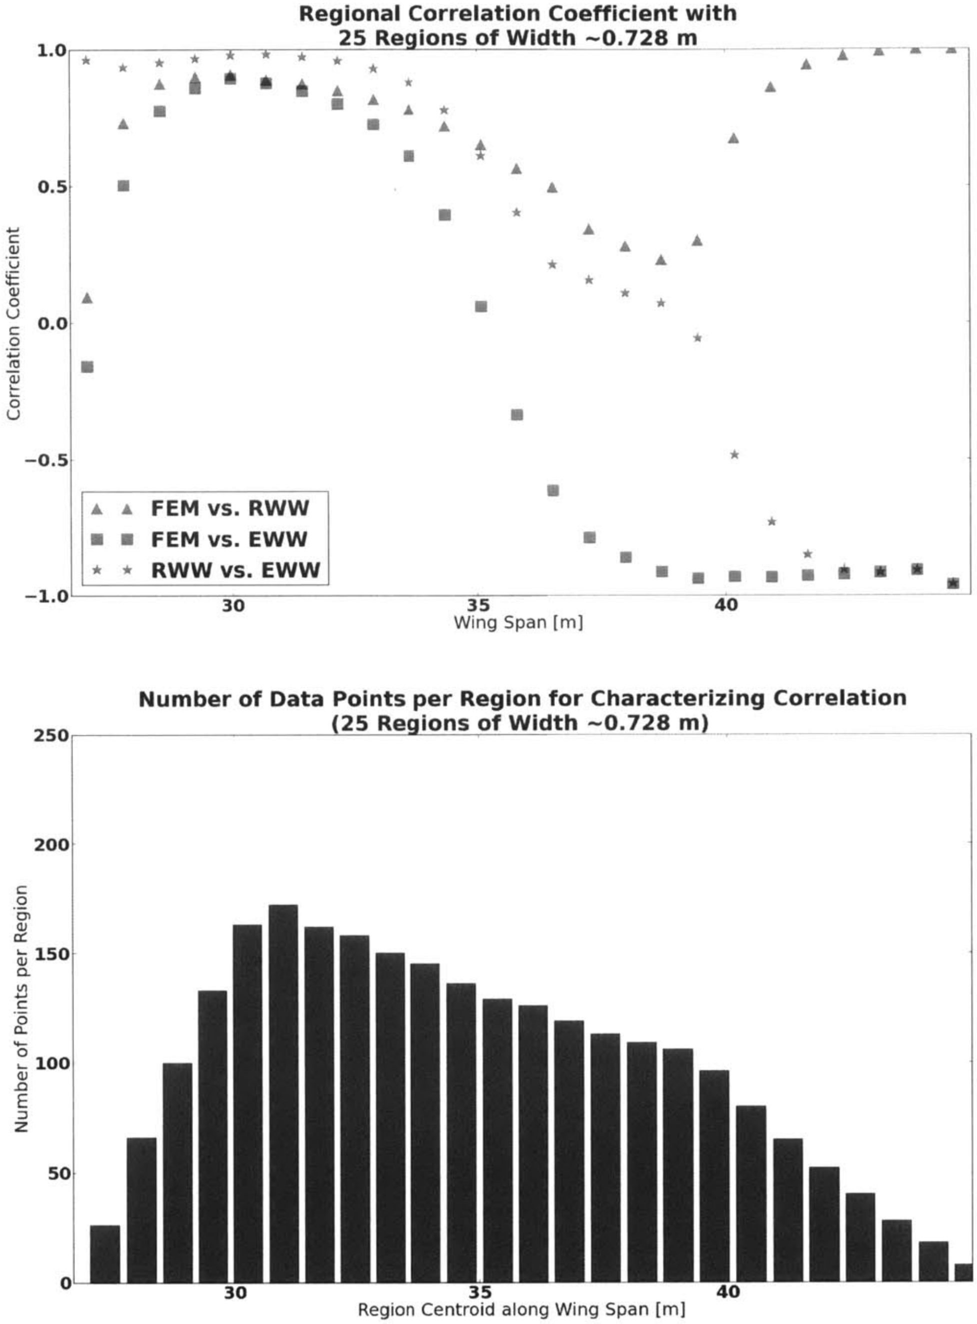

Figure 3-3: The design space is divided into 25 equally-sized regions. Top: the markers represent correlations between different pairs of models at region centroids. Triangles correspond to the correlation between the FEM and the RWW. Squares correspond to the correlation between the FEM and the EWW. Stars correspond to the correlation between the RWW and the EWW. Bottom: the bars correspond the the number of data points used for characterizing correlation. The left edge of each bar coincides with the centroid of its associated region.

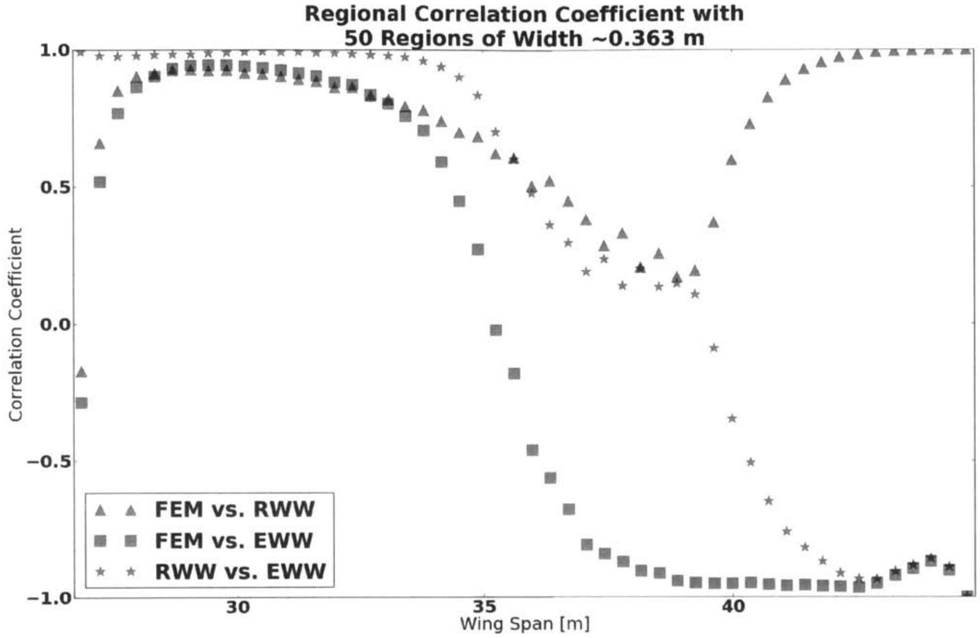

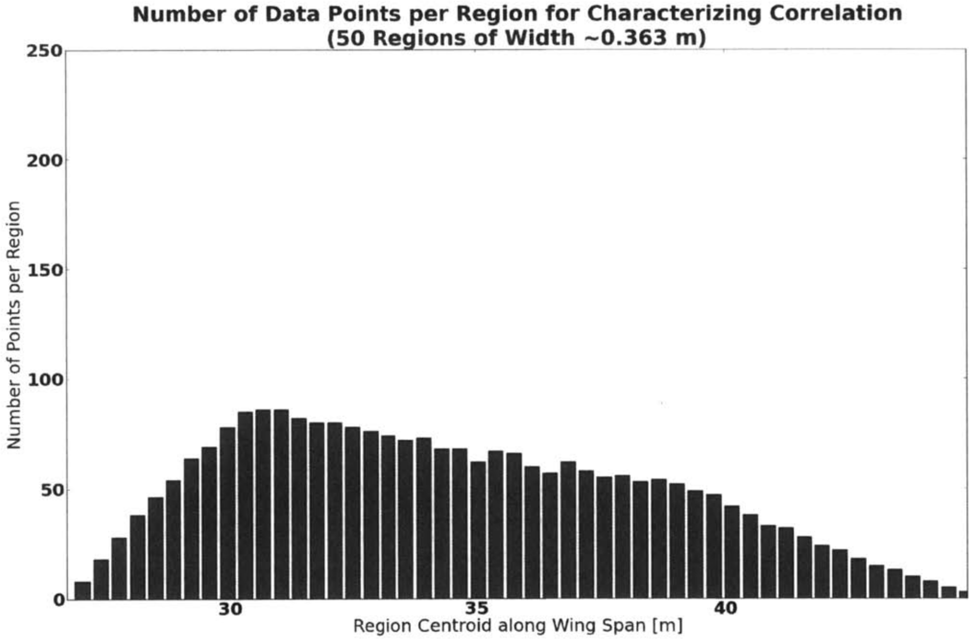

Figure 3-4: The design space is divided into 50 equally-sized regions. Top: the markers represent correlations between different pairs of models at region centroids. Triangles correspond to the correlation between the FEM and the RWW. Squares correspond to the correlation between the FEM and the EWW. Stars correspond to the correlation between the RWW and the EWW. Bottom: the bars correspond the the number of data points used for characterizing correlation. The left edge of each bar coincides with the centroid of its associated region.

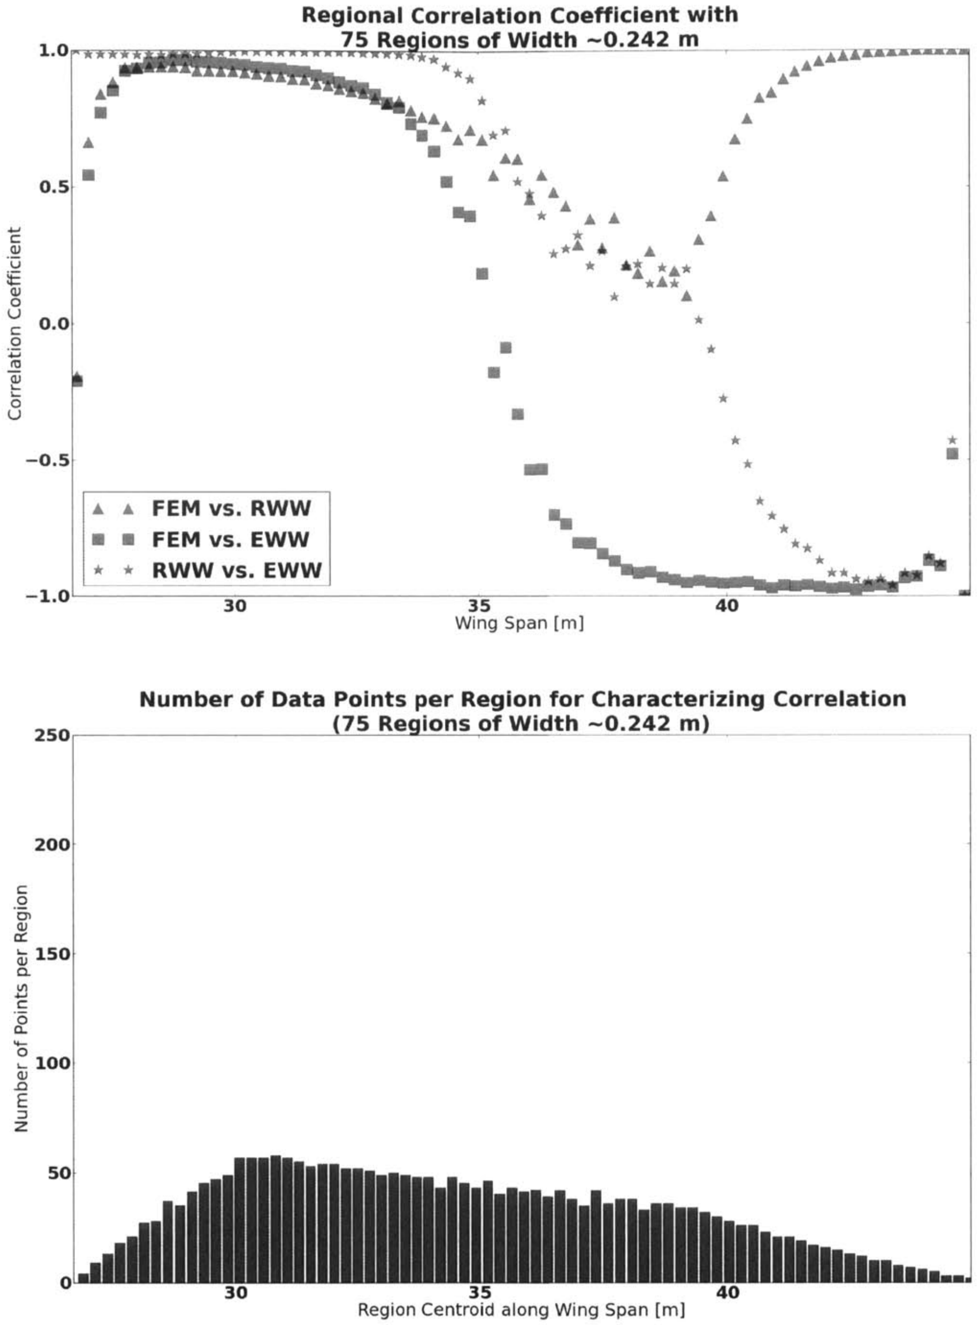

Figure 3-5: The design space is divided into 75 equally-sized regions. Top: the markers represent correlations between different pairs of models at region centroids. Triangles correspond to the correlation between the FEM and the RWW. Squares correspond to the correlation between the FEM and the EWW. Stars correspond to the correlation between the RWW and the EWW. Bottom: the bars correspond the the number of data points used for characterizing correlation. The left edge of each bar coincides with the centroid of its associated region.

with fewer data points carry larger degrees of uncertainty, since each data point holds considerable leverage in the collinearity of the discrepancy trends. This limitation can be alleviated by making the design space narrower or by obtaining additional data points. However, as we shall soon discuss, this is not always the most sensible solution. It is advisable to choose a stratification scheme that maintains a sufficient number of data points across all regions.

Another important fact to note is that the correlation averaged from two adjacent regions is not the same as the correlation from the union of the two regions. In particular, if the discrepancy trends for the models follow some nonlinear pattern closely, then the magnitude of correlation from the union of adjacent regions will generally be lower than that from each individual region. Therefore, it is possible for the model correlations to change drastically as the region widths vary. Perhaps a more fitting way to think about stratification then is to consider region widths, and let the number of regions follow naturally. As such, we can define regions that both contain a sufficient number of data points and are physically sensible. A wider region captures physics on a larger scale, while a narrower region captures physics on a finer scale. This idea is easily illustrated through drag aerodynamics, where higher order effects often ignored in low Mach regions can become significant towards high Mach regions. Hence, if Mach number was a design variable, the transition from low Mach number to high Mach number would prompt a corresponding shift in region width across the design space. As a result, we only need to constrict the design space or obtain new data points if there is still an insufficient number after appropriate region widths have been chosen. The execution of the proposed idea carries a high degree of subjectivity, and relies on proper judgment on behalf of the engineer. An inappropriately chosen region width may introduce unnecessary noise to the final surrogate or leave out valuable correlation information. Nevertheless, our observations in Figures 3-2, 3-3, 3-4, and 3-5 suggest that correlation trends are relatively robust against different stratification schemes, or at least situationally so, given the lack of a solid mathematical basis. As an additional verification, we simulate continuous correlation by imposing a "moving window" across the design space that expands and contracts dynamically to include a constant number of data points everywhere. Figure 3-6 shows the result. The general correlation trend from various stratification schemes matches that of the continuous simulation. Note that the computational cost for continuous simulation follows the well-known class of nearest neighbor search algorithms, and can become prohibitively expensive as the number of dimensions increases.

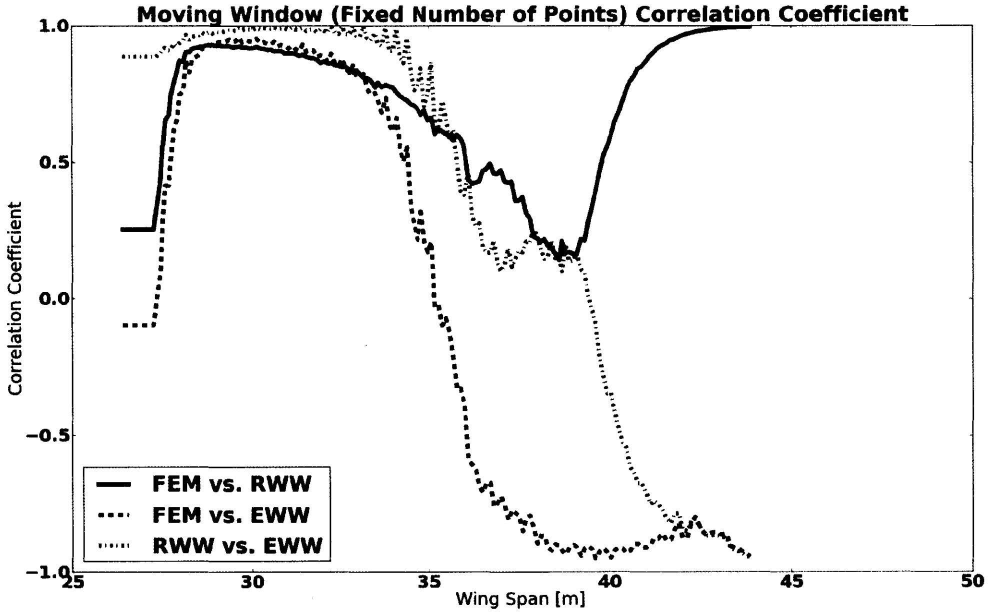

Figure 3-6: The black lines represent model correlations evaluated at 200 evenly-spaced wing span values using 50 data points per evaluation.

For our problem, we set the region width to be uniformly around 1.0 m, which gives us a sufficient number of data points across all regions. Moreover, the region width of 1.0 m agrees with the order of magnitude on which the model physics act. The significance behind the region width is similar to that of the characteristic length from the squared exponential covariance function used to create the GP surrogates of our models. In Chapter 4, we will address the relevant hyperparameters in greater detail. Based on experience, a reasonable region width usually lies between 0.01 and 0.1 times the GP surrogate's optimal characteristic length.

At present time, our framework does not quantify and propagate uncertainty associated with model correlation into the final surrogate. This is an important factor to consider both when applying our ideas and for future work. The aforementioned uncertainty stems from uncertainty in region widths, and inconsistent distribution of data points per region.

3.1.4 Interpolating for Continuity

The information fusion process operates by taking a correlation adjusted, variance-weighted average over constituent surrogate means. Consequently, the continuity of the fused result is dependent upon that of the individual surrogate means, variances, and model correlation. Any discontinuity in constituents will translate to discontinuity in the fused surrogate. While dividing the design space into regions allows us to capture variability in model correlation, we still face the problem of enforcing continuity between regions. The continuous simulation from previous section may seem promising, but can become prohibitively expensive in higher dimensions. It also carries added complexity from simultaneously handling the "number of points per region" and the "region width" constraints (in essence, it solves a constrained nearest neighbor search problem every time a local correlation is queried). A solution then is to interpolate between the centroids of adjacent regions. However, not all interpolation methods can be applied. An additional property we must preserve is the positive-definiteness of the correlation matrix $\boldsymbol{\rho}$, represented in the three-model case by

$$\boldsymbol{\rho} = \begin{bmatrix} 1 & \rho_{A,B} & \rho_{A,C} \\ \rho_{A,B} & 1 & \rho_{B,C} \\ \rho_{A,C} & \rho_{B,C} & 1 \end{bmatrix}$$

where $\rho_{A,B}$ represents the correlation between models A and B, and so on. This property guarantees the proper conditioning of the covariance matrix $\boldsymbol{\Sigma}$ used in the fusion formula, as well as ensures the final surrogate variance is greater than zero, thus physically sensible. If $\boldsymbol{\rho}$ is not positive-definite everywhere and model correlations are continuous, then at some point an eigenvalue of the covariance matrix will drop to zero, leading to numerical instabilities.

To find an interpolation scheme that preserves positive-definiteness, we use lemma 2.2 in Fitzgerald[12], which states that any positive linear combinations of real positive-semidefinite matrices is also positive-semidefinite. In application to our problem, we perform an Euclidean-distance-weighted linear interpolation of the correlation matrices from each pair of adjacent regions. Figure 3-7 shows the result. The case where ρ possesses a zero eigenvalue only occurs when a pair of models have correlation of 1—a situation only possible when all the data points within a certain region are perfectly aligned. This would signify a problem with model physics or region selection rather than a by-product from interpolation. Therefore, ρ should never contain a zero eigenvalue, and the method guarantees positive-definiteness.

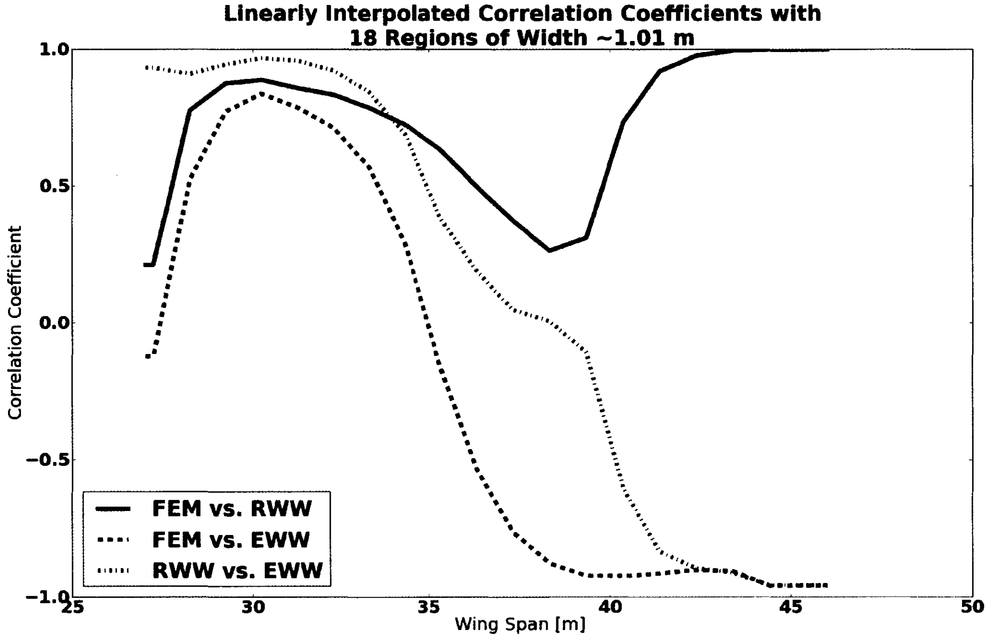

Figure 3-7: The solid line represents the interpolated model correlation between the FEM and the RWW. The dashed line represents the interpolated model correlation between the RWW and the EWW. The dotted line represents the interpolated model correlation between the RWW and the EWW.

While Figure 3-7 demonstrates the idea in 1-D, it can easily be extended to multiple dimensions. A drawback to this method is that while continuity can be guaranteed in the final surrogate, smoothness cannot. The result of information fusion can only be as smooth as its constituent means, variances, and correlations. In order to conduct gradient-based optimization on the final surrogate, post-smoothing would need to be applied.

3.1.5 Relation to Model Physics

We discussed earlier that model correlation refers to the degree of linearity in model discrepancy trends. Similar discrepancy trends do not necessarily imply same mechanisms behind the models. Therefore, we provide explanations of the linear interpolated correlation trends based on our knowledge of model physics. Figures 3-8, 3-9, and 3-10 show how model discrepancies match up with correlation trends in respective model pairs. The plots divide the correlations between different pairs of models into three colored boxes. Red box corresponds to low wing spans (~25-34 m), yellow box corresponds to medium wing spans (~34-40 m), and blue box corresponds to high wing spans (~40-46 m). For the FEM vs. the EWW, the plot of discrepancies along wing span may appear misleading, so we refer to Figure 3-11 to aid visual inspection.

Starting with the FEM vs. the RWW, we notice that the two are highly correlated in the red box. That is, taking slices along wing span (z-axis), the slope of the linear fit between the FEM and the RWW discrepancies is positive, and the residuals are small. In this box, the RWW's own wing sizing routines converge on wing geometry that is very similar to what is fed directly into the FEM, so the high positive correlation is indeed due to similar model physics. Correlation between the two models falls sharply in the yellow box, which is due to the linear fit between the model discrepancies worsening. This happens because in the yellow box, the RWW is more sensitive to variables aside from wing span, and its wing sizing routines produce corrections that are absent in the FEM over the associated range of inputs. In the yellow box, the RWW's discrepancies are somewhat more of a constant offset than scaled on inputs in comparison to the FEM's. The correlation is still positive because the slope of the linear fit is still positive. Finally, the model correlations bounce back up in the blue box. This is because the RWW assumes a fully exposed (no fuselage diameter) rectangular wing, which, at high wing spans, carry many of the same properties as the simple hollow rectangular beam used by the FEM. Thus, despite the differences in the mean offset of the two—a consequence of misrepresenting inboard and outboard geometries by different degrees, the models operate under similar physics.

Moving onto the FEM vs. the EWW, we notice that the models are highly correlated in the red box. This can be explained simply by the observation that the elliptic wing geometry and lift profile are not too different from the rectangular counterparts assumed by the FEM at low wing spans. In the yellow box, model correlations transition from positive to negative. Taking slices along wing span, we would see the slope of the linear regression line go from slightly positive to slightly negative. The correlation between the FEM and the EWW decreases sharply for the same reason as in the case for the FEM vs. the RWW (the EWW is sensitive to variables other than wing span). The correlation does not bounce back up because as wing span increases, the overcompensation of the EWW's fuselage carryover becomes significant, causing the EWW to overestimate wing weight while the FEM continues to underestimate it. The differences in physics remain meaningful at high wing spans, so it is expected that the two models exhibit medium-high negative correlation in the blue box.

Next, we examine the RWW vs. the EWW. The models are positively correlated in the red box. The high correlation is the result of both models having an internal sizing routine, thus computing wing box geometry in a similar manner (e.g. spar cap thickness based on bending moment). Furthermore, aside from several fixed parameters unique to the EWW for estimating fuselage carryover lift (e.g. fuselage diameter and wing attachment location), both models use the same set of inputs to size the wing. In the yellow box, model correlation steadily decreases. This follows from the differences in wing geometry, lift profile, and fuselage carryover between the models gradually becoming significant. The EWW's discrepancies initially trend downwards with increasing wing span, but turns to trend upwards once its overcompensation for fuselage carryover becomes problematic. The RWW's discrepancies always trend downwards with increasing wing span. This explains the eventual sign switch in correlation in the yellow box. At high wingspans, the EWW's physics continues to push its discrepancies upwards, while the RWW's discrepancies continue to trend in the opposite direction. As a result, the models are negatively correlated in the blue box.

The three colored boxes cover the entire design space of about 21.0 m with negligible amounts of overlap. Given our choice of region width at approximately 1.0 m, each box should contain multiple regions. The explanations for correlation trends therefore only apply macroscopically. In general, for some region centroid c_i, we can project all points in that region onto the plane at $z = c_i$ in order to study the local model correlation. Physical explanations of observed trends can be derived by studying the overall similarities and differences among the models and identifying the most dominant ones in the region of interest.

3.2 Fidelity Variance Assignment

This section covers the method for assigning variance associated with model fidelity in the context of our demonstration problem.

3.2.1 Polynomial Regression

In surrogate based optimization, there are two main sources of uncertainty. The first is uncertainty due to sparsity of samples in the training set, and the second is uncertainty due to model fidelity. Since the models used in engineering design are almost always deterministic, there is no uncertainty due to measurement noise. Uncertainty due to sparsity of samples can be quantified through Bayesian regression. Oftentimes, this source of uncertainty can be mitigated if prior knowledge regarding underlying structure of the data is available. If we have no information aside from simple governing physics behind the models, it is generally good practice to use kernel methods. Uncertainty due to model fidelity is, however, far more subjective. There is no standard way to characterize model fidelity, and the expert elicitation process that precedes it is a field of its own. Given that the focus of our project is on utilizing model correlation to aid information fusion, we only provide a loose guideline for the assignment of fidelity variance to respective models.

The data used for characterizing correlation is in the form of discrepancy from TASOPT. Thus, in addition to studying discrepancy trends between different model pairs, we can use the data to estimate the performance of the models across the entire design space. Once again, we emphasize that the coverage range of the data should fully encompass the design space. Otherwise, we would have to extrapolate both model correlation and model fidelity variance, making the surrogate highly unreliable. Since each model performs differently at different locations, one way to represent model fidelity is to fit a curve through the discrepancies of each model and treat the absolute value of the response surface as some standard deviation. The particular polynomial used for the curve fit in the demonstration problem is quadratic, and is chosen based on visual inspection of the data. The variance is then simply the square of that standard deviation.

3.2.2 Addition of Constant Term

In addition to the polynomial fit, we add a constant term on the order of the maximum absolute discrepancy of each model to their respective fidelity standard deviations. This way, the true wing weight is more likely to fall within the confidence interval associated with our model estimates. The magnitude of the constant is again at the engineer's discretion. For the demonstration problem, we assign 1500 kg to the FEM. This is a conservative estimate, but one we adopt to demonstrate a few properties of fusion with correlation later on. For the RWW, we assign 100 kg. The small value is justified because despite misrepresenting the wing in a variety of ways, the RWW still performs well in general. Finally, we assign 500 kg to the EWW. The value is chosen because the EWW overcompensates for fuselage carryover by failing to reduce the structural requirements of the in-fuselage portion of the wing, the magnitude of which averages to around 500 kg. Figure 3-12 shows the result after the constants have been added. Note that the constants are not added to standard deviations directly, but are squared and then added to the respective variances. This way we introduce a smoothing effect to the standard deviations.

Given that the posterior distribution for our choice of Bayesian regression should be analytically tractable (otherwise the added computational cost defeats the purpose of surrogate optimization), we prioritize convenience over precision. Thus, despite the fact that certain models may only underestimate or overestimate the true output, our surrogates employ symmetric posterior densities that cannot accommodate one-sided confidence intervals.

Despite what Figure 3-12 may suggest, the method we used to assign fidelity variance does not require the models to follow some strict hierarchy. It is also worth mentioning that regardless of how one chooses to characterize model fidelity, continuity in fidelity variance functions should be enforced. Otherwise, the fused surrogate will also display discontinuities at the corresponding locations.

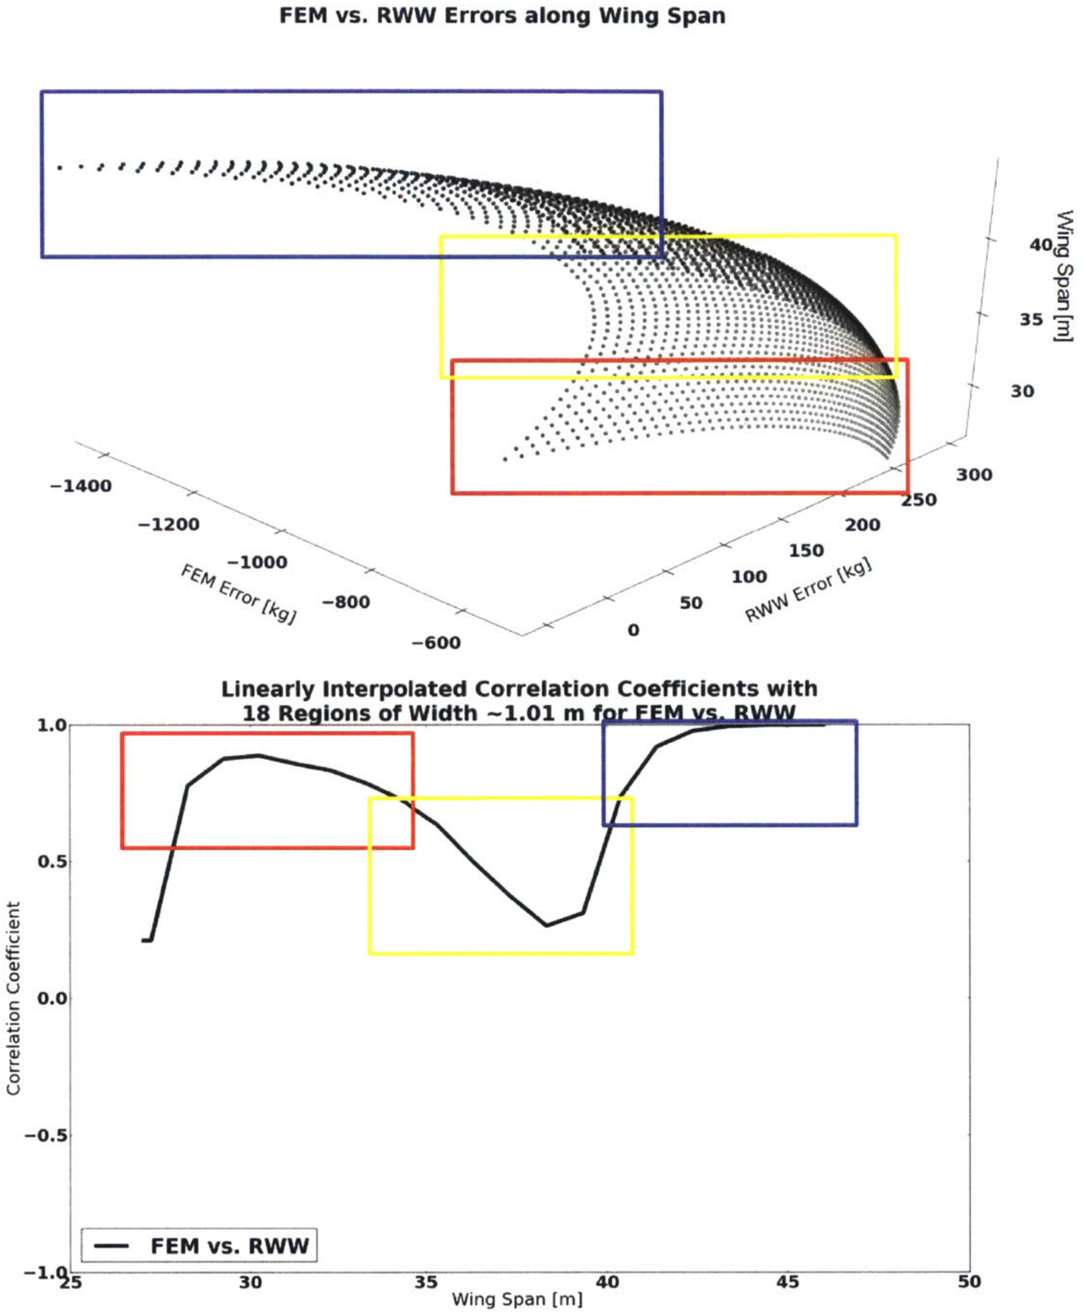

Figure 3-8: Top: the points represent model discrepancies of the FEM and the RWW along wing span. Bottom: the solid line shows the model correlation between the FEM and the RWW computed based on their discrepancy trends. The red box encompasses low wing spans, the yellow box encompasses medium wing spans, and the blue box encompasses high wing spans, classified with respect to the range of the design space.

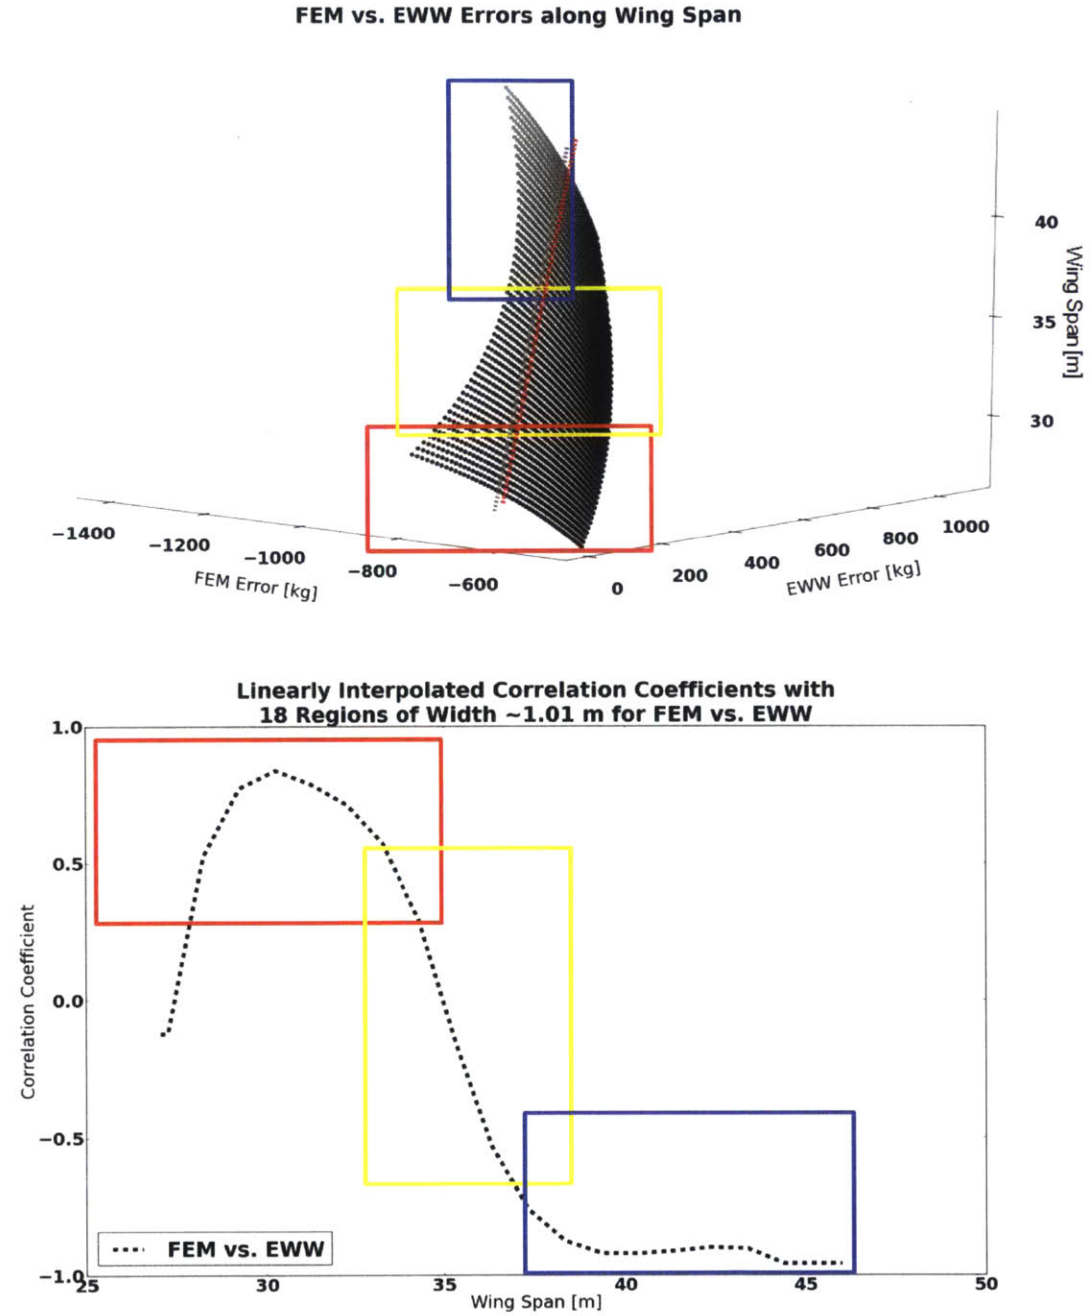

Figure 3-9: Top: the points represent model discrepancies of the FEM and the EWW along wing span. Bottom: the solid line shows the model correlation between the FEM and the EWW computed based on their discrepancy trends. The red box encompasses low wing spans, the yellow box encompasses medium wing spans, and the blue box encompasses high wing spans, classified with respect to the range of the design space.

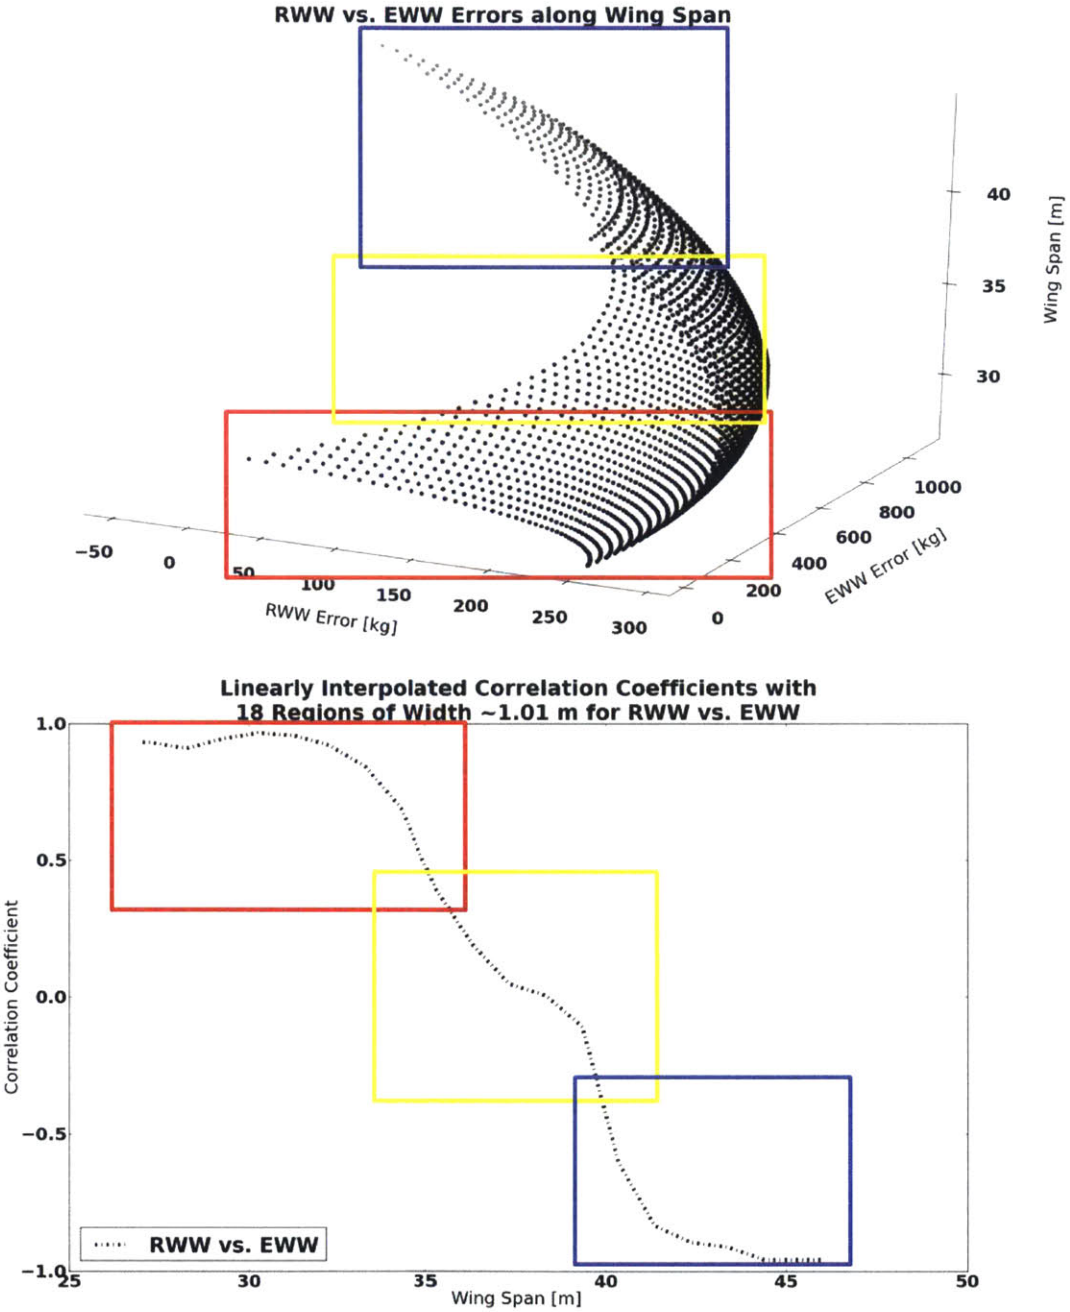

Figure 3-10: Top: the points represent model discrepancies of the RWW and the EWW along wing span. Bottom: the solid line shows the model correlation between the RWW and the EWW computed based on their discrepancy trends. The red box encompasses low wing spans, the yellow box encompasses medium wing spans, and the blue box encompasses high wing spans, classified with respect to the range of the design space.

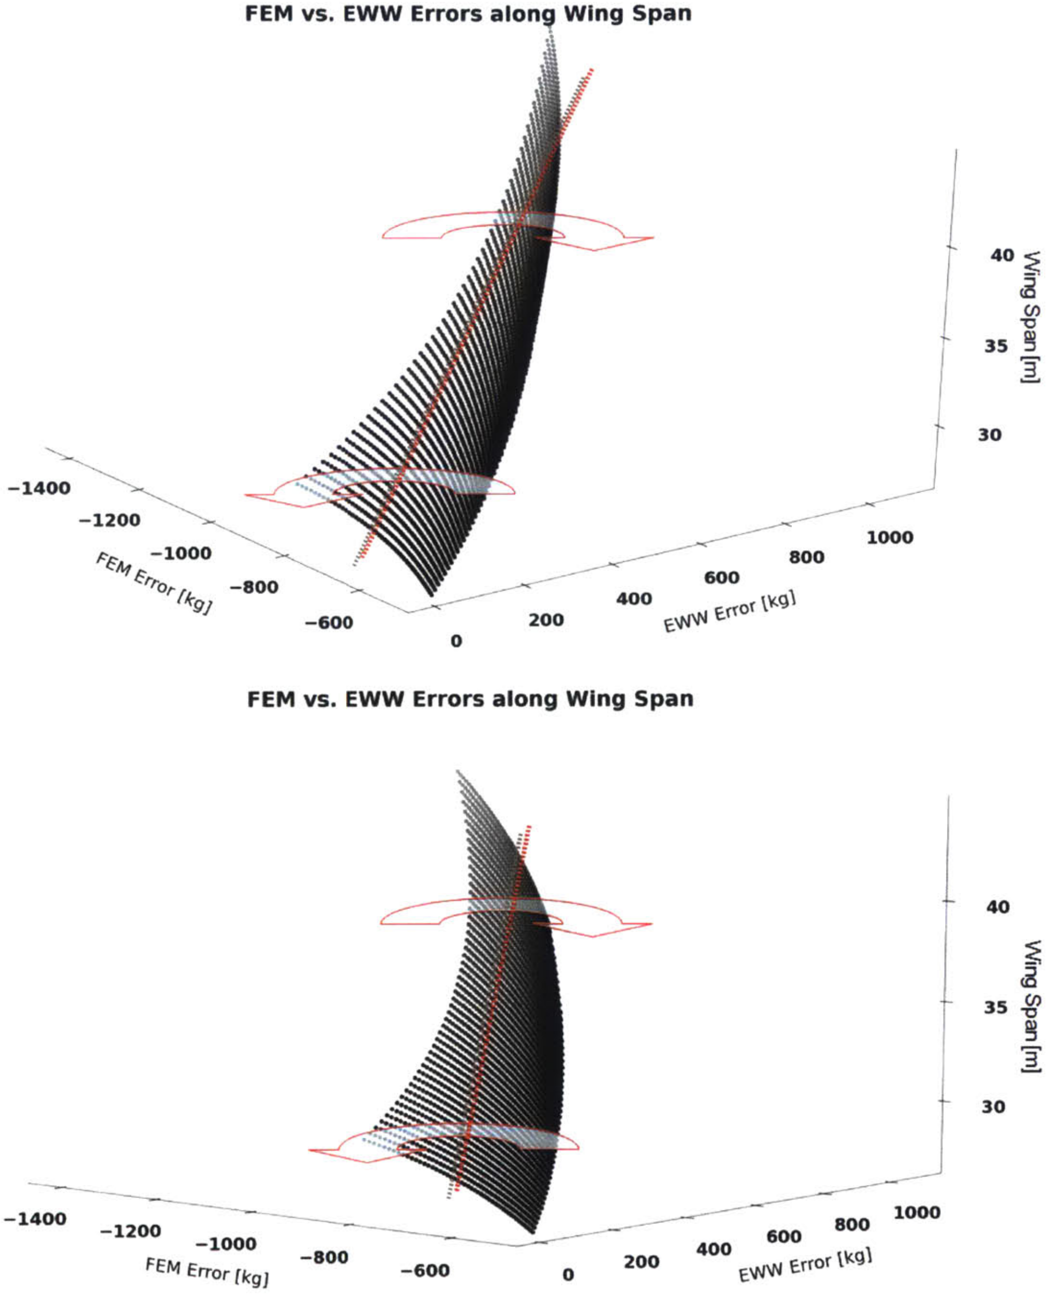

Figure 3-11: The points show the FEM vs. the EWW model discrepancies from different viewing angles.

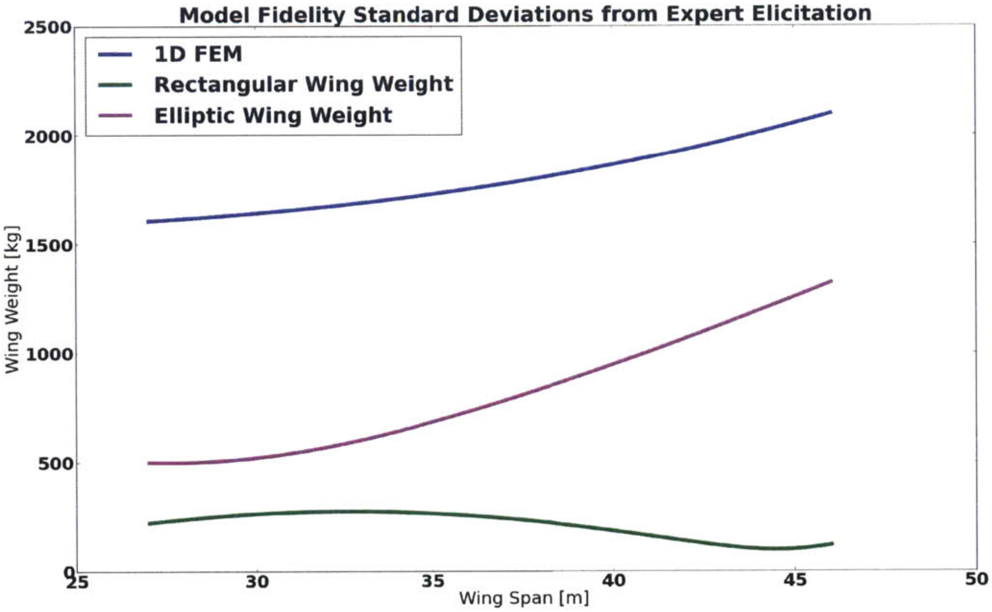

Figure 3-12: The colored lines represent 1 standard deviation from respective model output wing weight estimations. The probability of the true wing weight as approximated by each model follows a Gaussian distribution. The blue line represents the std. of the FEM, the green line represents the std. of the RWW, and the purple line represents the std. of the EWW, as functions of position in the design space.

Chapter 4

Fusion of Information

In this chapter, we apply regression techniques to generate surrogates for each model, and conduct fusion of correlated information based on those surrogates. Section 4.1 implements Bayesian regression on select set of training points for each model. Section 4.2 fuses the constituent surrogates, taking into account model correlation.

4.1 Individual Model Surrogates

This section introduces the individual model surrogates produced using Gaussian Process Regression for the demonstration problem.

4.1.1 Training Set

Unlike the application of kernel methods in machine learning and statistics, surrogate optimization works with very few training points while relying heavily on appropriately chosen kernel functions and parameter bounds to maintain the integrity of the results and minimize computational cost. To emulate the sparsity of training points typically available in design optimization problems, we restrict ourselves to just a few points in each training set. Furthermore, the training points are not distributed evenly across the design space, but are more concentrated in places where the respective models perform best. Figure 4-1 details the distribution of training points for each model. Note that the inputs used to produce the training points are derived from the same reference design as the data for characterizing model correlation and fidelity.

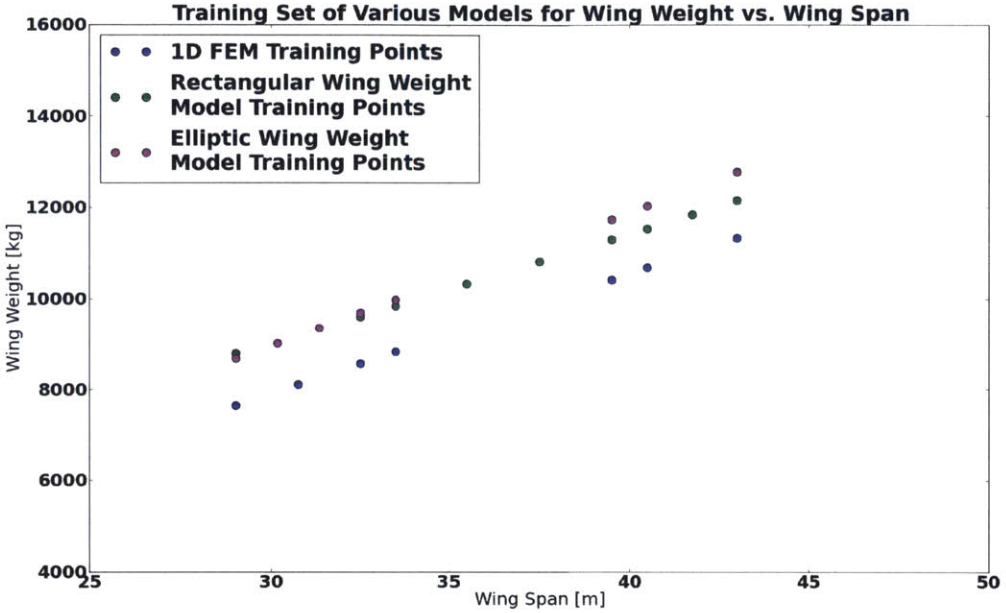

Figure 4-1: The colored points represent the training set for individual surrogates. Blue corresponds to the training set for the FEM, green corresponds to the training set for the RWW, and purple corresponds to the training set for the EWW.

It is generally good practice to include at least a few training points towards the edges of the design space regardless of local model fidelity. The reason is that many popular kernels used in building surrogates compute correlation matrices by some measure of statistical distance (e.g. Mahalanobis, Euclidean, etc.). When calculating the optimal parameters in these kernels through maximum likelihood, having training points with joint coverage over a large portion of the design space will generally lead to better estimates. In particular, since we apply Gaussian Process Regression with squared exponential kernel in the demonstration problem, having points towards the edges of the design space will aid the convergence onto an appropriate characteristic length scale.

4.1.2 Bounds on Hyperparameters

Gaussian Process Regression accommodates a broad range of kernel functions. Prior knowledge regarding the nature of the response surface should serve as a guideline for the its selection. Since all the major factors influencing aircraft wing sizing vary continuously, the square exponential kernel, which guarantees smoothness in the response surface, is appropriate. Equation 4.1 gives the exact form of the kernel in m-dimensions.

$$k(\boldsymbol{d^k}, \boldsymbol{d^j}) = \sigma_{GP}^2 \, exp(-\sum_{i=1}^{m} \frac{(d_i^k - d_i^j)^2}{2l_i^2}) \tag{4.1}$$

where $\boldsymbol{d^k}$ and $\boldsymbol{d^j}$ represent two training points, i represents the particular dimension in $\boldsymbol{d}$, σ_{GP}^2 represents variance due to sampling, and l_i represents the characteristic length scale in the i^{th} dimension.

The hyperparameters of interest are σ_{GP}^2 and l_i. σ_{GP}^2 sets the upper bound on variance in the absence of nearby training points, and l_i can be understood qualitatively as the average distance one would need to travel on the response surface along dimension i before any significant change in profile is observed. While setting the hyperparameters by maximum likelihood right away can be enticing, we must first impose certain constraints derived from our knowledge of model physics so that the surrogates will be physically sensible. In general, bounds on the hyperparameters do not have to be the same across all models, but we keep ours consistent for convenience sake. First, we fix σ_{GP}^2 at 10000 kg^2 for all three models. σ_{GP}^2 should not be chosen automatically because it reflects the highest degree of potential errors when there are no training points nearby. Global maximum likelihood estimate makes use of patterns in the training set to infer information on the data-generating process. For σ_{GP}^2 in particular, the better the training points are aligned in the manner described by the response surface, the higher the confidence (the lower the σ_{GP}^2) that will be assigned to the regression estimates. This is a purely mathematical notion, and the resulting σ_{GP}^2 is in some ways decoupled from the model physics. A global maximum likelihood estimate would be justified if the data-generating process is susceptible to noise, or if the true mechanism that turns input to output is extremely complex and not well-understood (as is often the case in statistics). However, all of our models are deterministic and physics based, so σ_{GP}^2 should be chosen to be a conservative estimate that reflects the uncertainty from replacing the model output with some regressed approximation when there are no neighboring training points to serve as a guide. On the other hand, we allow some variability in l_i, given that no universal standard exists in determining the length scale on which the model physics act. In general, there can be multiple maxima in the likelihood function, but the global maximum is not always the most sensible one in the context of the design problem. The bounds we choose for l_i should ensure that the value found through local maximum likelihood remains physically sensible. Since we know beforehand that the relationship between wing span and wing weight is roughly linear, we expect the associated characteristic length scale to be rather large (there is almost no change in profile across any distance for a linear surface). Hence, we assign only a lower bound to l_i at 20 m for all three models. Note that the characteristic length scale converged upon by local maximum likelihood can often serve as a guideline for the selection of region width in characterizing model correlation, as these values now carry the appropriate physical significance.

4.1.3 Results

The result of the Gaussian Process Regression for respective models is shown in Figure 4-2. Note that the "measurement noise" is a regularization term added to the diagonal of the covariance matrix to ameliorate potential overfitting problems. The larger the magnitude of the term, the less stringent the regression is in getting the response surface to go through training points exactly. Dotted lines around the means represent one standard deviation after fidelity variance has been added. In this case, model fidelity appears to be the primary contributor to uncertainty everywhere in the surrogates, although this may not be true in general.

Table 5.1 summarizes the associated hyperparameters of the surrogates. As expected, the sampling variance for all three models is fixed at 10000 kg. The characteristic lengths all exceed the range of our design space ($\sim$ 21.0 m), which indicates

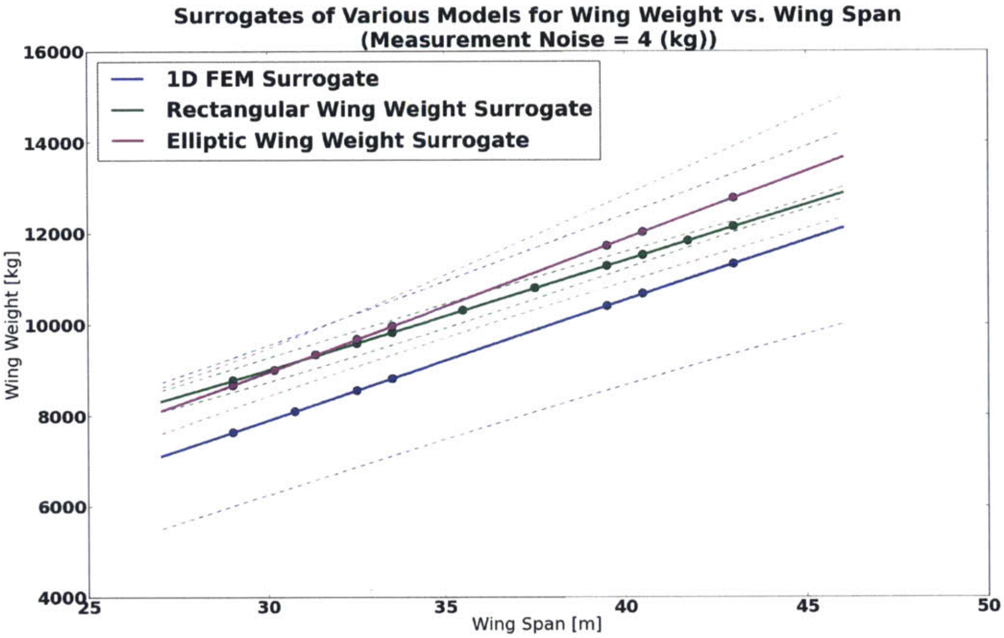

Figure 4-2: The solid lines represent the means of respective surrogates, and the dashed lines represent ±1 standard deviation about the means. Blue corresponds to the surrogate of the FEM, green corresponds to the surrogate of the RWW, and purple corresponds to the surrogate of the EWW.

	Sampling Variance σ_{GP}^2 $[kg^2]$	Characteristic Length l $[m]$
The FEM	10000	83
The RWW	10000	67
The EWW	10000	67

Table 4.1: The columns show the respective Gaussian Process Regression hyperparameters for the 1-D demonstration problem.

that the response surfaces are linear.

4.2 Information Fusion

This section details the algorithm for combining the individual surrogates under the presence of model correlation.

4.2.1 Theory and Review

Fusion of information without any restrictions on the form of participating probability densities is very difficult. Hence, the algorithm we present applies only to multivariate Gaussians. Given that model errors often assume normal distributions, and many state-of-the-art surrogate modeling techniques generate Gaussian posteriors, the algorithm is still valid across a wide range of setups. Our demonstration problem falls within this category, so the aforementioned conditions are satisfied. In addition, the algorithm does not require the surrogates to be produced under the same regression technique. This property allows the engineer to choose the most suitable regression method for each surrogate and still be able to fuse information effectively. Equation 4.2 details the posterior density of the true unknown output, along with its estimators,

$$\begin{aligned} p(z^*|\{Z_A|\boldsymbol{d}, Z_B|\boldsymbol{d}, ..., Z_{final}|\boldsymbol{d}\}) =& \frac{1}{\sqrt{2\pi \ var(Z^*)}} exp(-\frac{(z^* - \mathbb{E}[Z^*])^2}{2 \ var(Z^*)}) \\ \mathbb{E}[Z^*] =& \frac{\mathbf{e}^T\Sigma^{-1}\mathbf{z}_{1:t}}{\mathbf{e}^T\Sigma^{-1}\mathbf{e}} \\ var[Z^*] =& \frac{1}{\mathbf{e}^T\Sigma^{-1}\mathbf{e}}, \end{aligned} \tag{4.2}$$

where z^* is the true unknown output, Z^* is a random estimate of z^* based on information fusion, and $Z_A|\boldsymbol{d}, Z_B|\boldsymbol{d}, ..., Z_{final}$ are random estimates of the true output based on constituent surrogates under some common input $\boldsymbol{d}$. $\mathbf{e} = [1, ..., 1]^T$, $\mathbf{z}_{1:t} = [z_A|\boldsymbol{d}, z_B|\boldsymbol{d}, ..., z_{final}|\boldsymbol{d}]$, and Σ is the covariance matrix. For three constituent models, as is the case for our demonstration problem, the covariance matrix can be represented as

$$\Sigma = \begin{bmatrix} \sigma_A^2 & \rho_{A,B}\sigma_A\sigma_B & \rho_{A,C}\sigma_A\sigma_C \\ \rho_{B,A}\sigma_B\sigma_A & \sigma_B^2 & \rho_{B,C}\sigma_B\sigma_C \\ \rho_{C,A}\sigma_C\sigma_A & \rho_{C,B}\sigma_C\sigma_B & \sigma_C^2 \end{bmatrix}$$

which can be easily extended to N models. Specifically, if we enumerate the models with subscript $\mathcal{M}_i$, a generalized Σ can be written as

$$\Sigma = \begin{bmatrix} \sigma^2_{\mathcal{M}_1} & \rho_{\mathcal{M}_1,\mathcal{M}_2}\sigma_{\mathcal{M}_1}\sigma_{\mathcal{M}_2} & \cdots & \rho_{\mathcal{M}_1,\mathcal{M}_N}\sigma_{\mathcal{M}_1}\sigma_{\mathcal{M}_N} \\ \rho_{\mathcal{M}_2,\mathcal{M}_1}\sigma_{\mathcal{M}_2}\sigma_{\mathcal{M}_1} & \sigma^2_{\mathcal{M}_2} & \cdots & \rho_{\mathcal{M}_2,\mathcal{M}_N}\sigma_{\mathcal{M}_2}\sigma_{\mathcal{M}_N} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{\mathcal{M}_N,\mathcal{M}_1}\sigma_{\mathcal{M}_N}\sigma_{\mathcal{M}_1} & \rho_{\mathcal{M}_N,\mathcal{M}_2}\sigma_{\mathcal{M}_N}\sigma_{\mathcal{M}_2} & \cdots & \sigma^2_{\mathcal{M}_N} \end{bmatrix}$$

Note that Σ is positive-definite and symmetric. The covariance matrix can be constructed by taking the entrywise product of the correlation matrix $\boldsymbol{\rho}$ and the matrix consisting only of the covariances among model discrepancies. $\sigma_{\mathcal{M}i}$ is equivalent to the corresponding surrogate's standard deviation after fidelity variance has been accounted for.

The algorithm for fusion of correlated information is supplied by Winkler[30], and demonstrated in the context of multifidelity design optimization by Allaire[3]. In his work, Allaire explored fusion of correlated information using two abstract models and a manually specified correlation coefficient. He made the observation that if two models are highly correlated and have very different fidelity levels, the fused result will exhibit even lower variance than if the models were independent to start with. He argues that this is because the adequacy of the lower fidelity model must have been understated (since it is so highly correlated with a model in which we have great confidence), and thus the lower fidelity model provides more information to the fusion than its level of fidelity implies. On the other hand, if the two models are highly correlated and have very similar fidelity levels, the fused result will exhibit higher variance than if the models were independent. This is more intuitive because the models' respective contribution of information to the fused result is not distinct, and redundant information should not be considered when evaluating confidence. Furthermore, the mean of the fused model will shift in the direction of the higher fidelity model as correlation increases, even to the extent of going beyond the original model's mean if the correlation is high enough. This can be explained by that highly correlated models are more likely to produce estimates that lie on the same side of the true quantity, since both models are prone to make the same mistakes.

In continuation of Allaire's work, Christensen explored fusion of correlated information in a multidisciplinary aircraft design problem[6]. Christensen found that by taking model correlation into account, the final design converged upon through surrogate optimization is slightly different from the independent case, and carries less confidence in its parameters. Hence, there is evidence that model correlation has the potential to make substantial impact on the outcome of surrogate optimization.

4.2.2 Independent Case

Fusion of information under the independence assumption has been demonstrated by Lam[22]. Mathematically, it is the same as using the identity as the correlation matrix $\boldsymbol{\rho}$. Figure 4-3 shows the result in application to our problem.

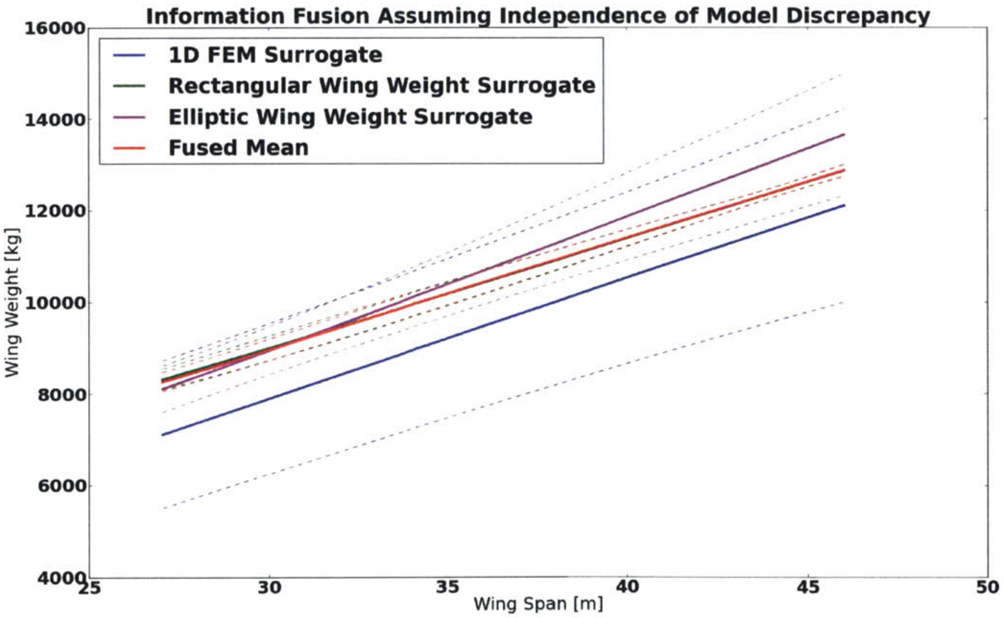

Figure 4-3: The solid red line corresponds to the mean of information fusion under model independence, and the dashed red curves symmetric about the mean correspond to ±1 standard deviation. The solid blue line corresponds to the mean of the FEM's surrogate, and the dashed blue lines correspond to ±1 standard deviation. In a similar manner, green lines correspond to the RWW's surrogate, and purple lines correspond to the EWW's surrogate.

Under the independence assumption, information from each constituent model does not overlap with information from any of the others. The fused mean is then a weighted average over all constituent means, and more weight is given to models with lower variances. Figure 4-3 reflects this quite well, as the fused mean almost perfectly aligns with that of the RWW, which has the lowest variance.

4.2.3 Correlated Case

When correlation is taken into account, the weight that fusion assigns each constituent model becomes modified by an additional factor. As suggested by Equation 4.2, the exact values of the weights can be computed by

$$\mathbf{w}^T = \frac{\mathbf{e}^T \Sigma^{-1}}{\mathbf{e}^T \Sigma^{-1} \mathbf{e}},$$

which, when multiplied to $\mathbf{z}_{1:t}$, yields the fused mean. Note that the weights sum to 1, but are not strictly positive. Negative weight occurs for model $\mathcal{M}_i$ if its correlations with other models are highly positive, and its variance is higher than that of the other models. Figure 4-4 shows the result of fusion with correlation for our problem.

As a consequence of the way we characterized model correlation, the fused mean is of differentiability class $\mathbf{C}^0$. Hence, before gradient-based optimization methods can be employed, post-smoothing must be performed in order to guarantee smoothness. One way to do so is to conduct another Gaussian Process Regression with small measurement noise on points densely sampled from the raw fused surrogate. As such, all the features of the fused surrogate can be preserved and the sharp edges smoothed without sacrificing confidence.

4.2.4 Analysis and Discussion

Upon initial inspection, there does not appear to be significant difference between the two fused surrogates. The fused mean for the correlated case lie slightly above that of the independent case, while the variance is generally lower. Figure 4-5 compares the two in greater detail.

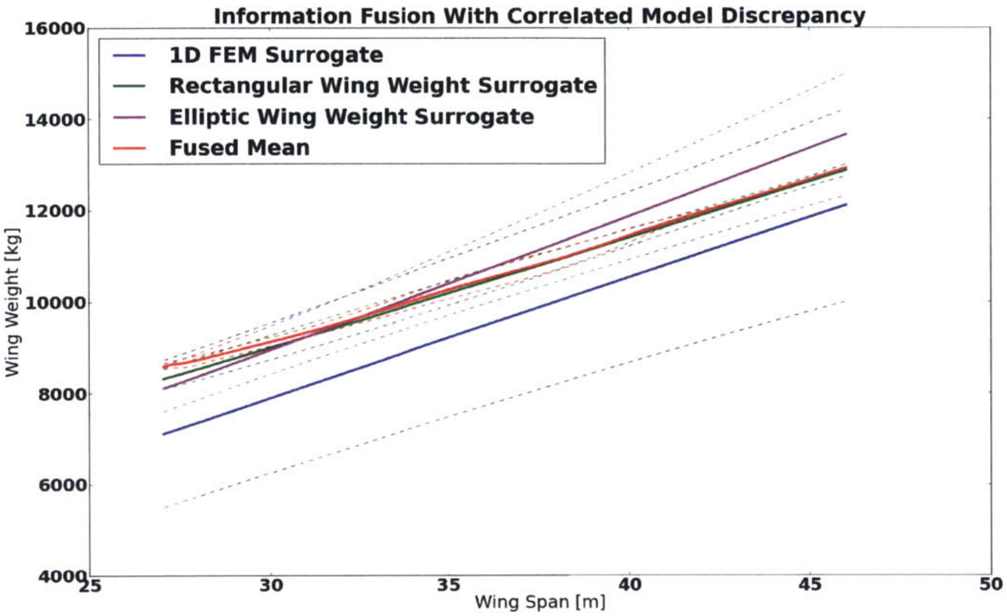

Figure 4-4: The solid red line corresponds to the mean of information fusion under model correlation, and the dashed red curves symmetric about the mean correspond to ±1 standard deviation. The solid blue line corresponds to the mean of the FEM's surrogate, and the dashed blue lines correspond to ±1 standard deviation. In a similar manner, green lines correspond to the RWW's surrogate, and purple lines correspond to the EWW's surrogate.

Since correlation and model uncertainty both vary continuously across the design space, it is not surprising that we see different behavior at different locations. At wing spans of about 27 to 34 m, fusion with correlation shows greater confidence than fusion under independence. This can be explained using the same argument Allaire made in that given the high positive correlation between the FEM and the other two models, the FEM's fidelity level must have been understated, and thus it provides more information to the fusion than its level of fidelity implies. At wing spans of about 35 to 40 m, fusion with correlation exhibits lower confidence than fusion under independence. This can be explained in that moderate correlation among the models suggests the presence of redundant information, which should lead to higher variance than if all information were distinct. At wing spans of about 40 to 46 m, fusion with correlation displays higher confidence than fusion under independence. This is once again in part due to the FEM's fidelity level being understated. However, another source of variance reduction can be traced to the negative correlations between the FEM and the EWW, and between the RWW and the EWW. Mathematically, negative correlation in certain model pairs makes corresponding elements in Σ negative, leading to a general inflation of the sum of all elements in the inverse of Σ (more elements in the adjoint of Σ are positive). Since Σ^{-1} lies in the denominator of the fused variance formula, the greater the sum of its elements, the lower the fused variance. The physical reasoning behind this phenomenon stands in that when models are negatively correlated, their discrepancy trends go in opposite directions, and inference on the direction of the true output can be made based on bounds set by the aforementioned discrepancy trends. Figure 4-6 tracks the denominator in the fusion formulae (i.e. $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$). As expected, the denominator increases sharply at wing spans of about 40 to 46 m, which is accompanied by an increase in correlation between the FEM and the RWW (the more positive this value, the larger the denominator), and an increase in negative correlation between the RWW and the EWW (the more negative this value, the larger the denominator).

In general, there are two main conflicting forces in characterizing variance. The first is variance reduction from extrapolated information based on model correlation (the higher the correlation, the more confidence the extrapolated information carries). The second is variance increase from inconsistency of information carried by highly correlated models. Interactions between model correlations and variances determine the degree to which each force manifests in the fusion.

Before the surrogate is ready for optimization, it is generally good practice to check the conditioning of various components. Even though we have guarantees of positive-definiteness in Σ, it does not preclude us from possible numerical instabilities derived as a result of interactions between terms during fusion. Figure 4-7 tracks the condition number of Σ and $\boldsymbol{\rho}$ across the design space.

Visual inspection reveals that the maximum condition number lies around 10^6. A condition number of 10^6 will lose us an upwards of 6 digits of accuracy. Given that real, non-integer variables used throughout our analysis are stored under double-precision floating-point format, this leaves us with 8 to 9 significant decimal digits.

Thus, we still have more than enough numerical precision after information fusion has been conducted. If, in a separate application of this framework, a matrix becomes ill-conditioned, one should first verify that Σ is positive-definite everywhere in the design space. Once that has been confirmed, the problem most likely lies in $\boldsymbol{\rho}$, which is influenced by model physics and region selection. The most common mistakes usually stem from inappropriately chosen region widths, which leads to insufficient data points in certain regions. Such practice will cause high linear dependency among columns of $\boldsymbol{\rho}$, which translates to ill-conditioned Σ.

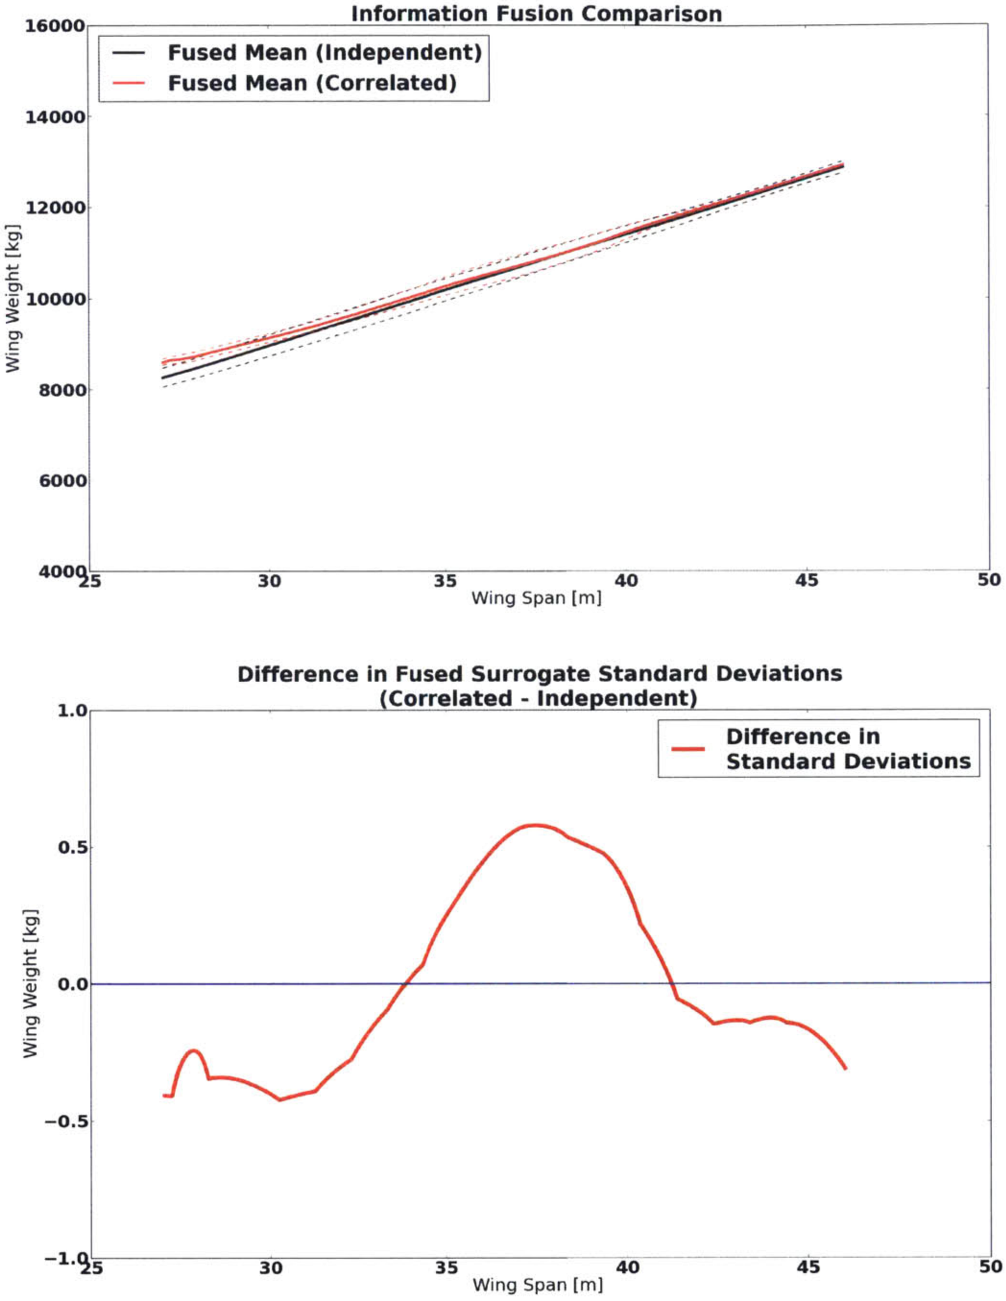

Figure 4-5: Top: the solid and dashed red curves correspond to the mean and ±1 standard deviation, respectively, of the fused surrogate under model correlation. The solid and dashed black lines correspond to the mean and ±1 standard deviation, respectively, of the fused surrogate under model independence. Bottom: the red curve represents the result after subtracting the standard deviation of the fused surrogate under model independence from the standard deviation of the fused surrogate under model correlation.

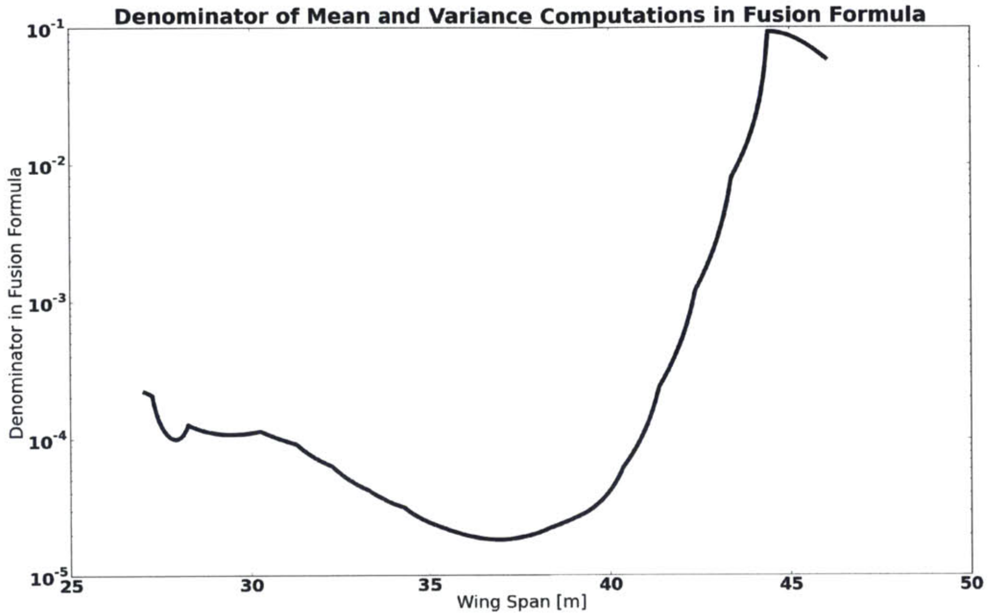

Figure 4-6: The black curve represents the quantity $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ (corresponding to the denominators of the mean and variance terms in Equation 4.2) across the design space. Notice that the y-axis uses logarithmic scale.

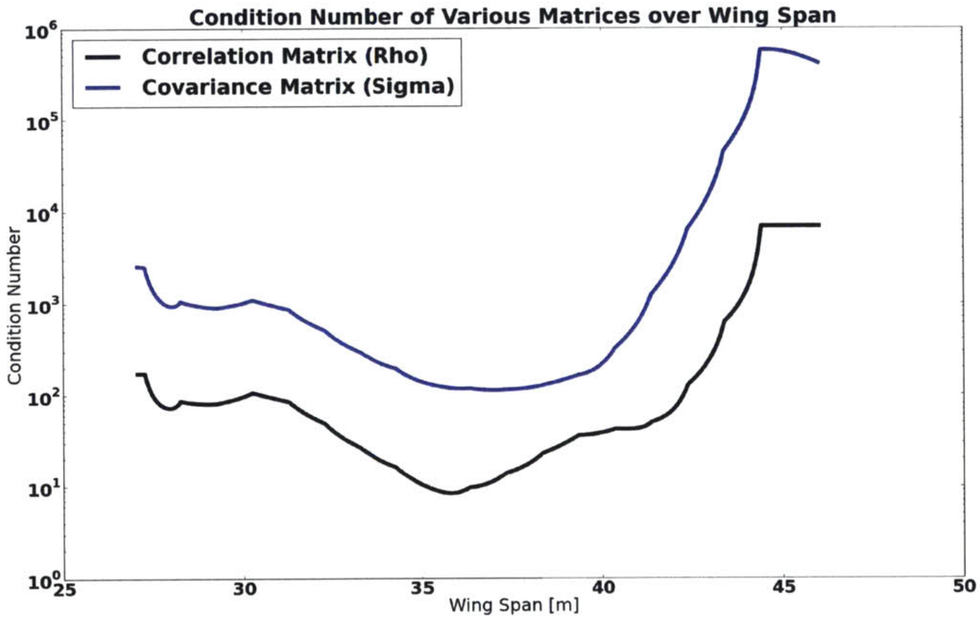

Figure 4-7: The blue curve represents the condition number of Σ across the design space. The black curve represents the condition number of $\boldsymbol{\rho}$ across the design space. Notice that the y-axis uses logarithmic scale.

Chapter 5

Application in 2-D

This chapter applies our proposed multifidelity framework on a 2-D problem. Section 5.1 outlines the setup of the surrogate optimization problem. Section 5.2 covers the characterization of model correlation. Section 5.3 overviews the assignment of fidelity variance. Section 5.4 presents the individual model surrogates built by different Bayesian regression methods. Finally, Section 5.5 shows the results from fusion of correlated information.

5.1 Problem Setup

As in the demonstration problem, we are interested in studying wing weight in relation to other design variables. We continue to focus our efforts on commercial transport aircraft, and keep the 737-800 as reference. Recall that in the demonstration problem, the chosen design variable was wing span. For the 2-D problem, we select aspect ratio and body weight (empty + payload - wings - engines weights). We choose aspect ratio because it is dimensionless. Dimensionless variables carry physical significance insensitive to scaling, so they can provide intuition over a wide range of wing architectures. We choose body weight because it does not specify wing geometry like aspect ratio, and instead impacts the wing weight by influencing wing sizing. Both variables are used only by the RWW and the EWW in computing wing weight. We intentionally exclude variables used by the FEM in order to demonstrate the versatility of our framework. Consequently, the output of the FEM should stay at the estimated wing weight of the 737-800 across the entire design space.

We want to explore wing weight for designs with aspect ratios spanning from 8.40 to 9.60 and body weights spanning from 34890 kg to 35600 kg. To generate data for characterizing model correlation and fidelity, the input to TASOPT must produce parameters encompassing the entire design space. Aspect ratio can be specified explicitly in the input, but body weight is computed by adding together a number of different component weights. Only certain component weights can be specified explicitly (e.g. passenger weight, etc.) in the input, while others are sized internally. Hence, we need one or more "tuning knobs" to control body weight, much in the same way we did for wing span in the demonstration problem. Fortunately, varying altitude alone is enough to influence body weight, and there is no need to find other highly correlated variables.

For generating the data, we sampled aspect ratios from 8.0 to 10.0 in 75 evenly-spaced segments, and altitudes from 29000 ft to 41000 ft in 75 evenly-spaced segments. The resulting body weights span from 34500 kg to 36000 kg, sufficiently covering our design space. There are 5625 unique designs in total, prior to the removal of any designs significantly different from our reference.

5.2 Characterizing Correlation

This section applies our method for characterizing model correlation to the 2-D surrogate optimization problem.

5.2.1 Data for Correlation

Out of the 5625 designs, a majority deviates considerably from the reference aircraft. In deciding which designs to keep for computing model correlation, we picked two parameters that are especially relevant in governing wing weight—wing span and engine weight. Wing span directly specifies wing geometry, while engine weight affects the degree of moment relief at wing root. For simplicity, we adopt a strict cutoff in relation to similarity of each design to the reference aircraft (737-800). Similarity is measured by Euclidean distance between parameter values. Given that the reference aircraft has a wing span of 34.18 m, we accept only designs with wing spans from 31 m to 37 m (approximately ±3 m from the reference). For engine weight, the reference aircraft has 5708 kg. We accept only designs with engine weights from 4700 kg to 6700 kg (approximately ±1000 kg from the reference). This idea can be extended to all other parameters, but since they either do not vary much in our data or are not especially important to wing sizing, we do not devise acceptance criteria with respect to them.

Once the data has been trimmed, we are left with 1889 unique designs. Figure 5-1 shows the errors of the FEM, the RWW, and the EWW based on those points projected onto the design space. As long as region divisions are properly made, 1889 designs should still be sufficient for characterizing correlation. It is worth mentioning that the data points are not distributed evenly across the design space, and we must take care to ensure that each region contains a sufficient number of points.

5.2.2 Region Selection

As in the demonstration problem, we select region widths based on model physics. For simplicity, we set region widths to be uniform across the design space. For body weight, we choose 740 kg, which leads to 2 regions along that dimension in total. 740 kg is appropriate because changes in output wing weight over our range of input body weight is small. A long region width captures general model correlation trends without being excessively sensitive to small variations. For aspect ratio, we choose 0.3, which leads to 5 regions along that dimension in total. 0.3 is appropriate because wing weight changes by roughly the same amount over 0.3 in aspect ratio as over 740 kg in wing weight. Notice that wing weight changes more over aspect ratio (as a fraction of its range in design space) than over body weight (as a fraction of its range in design space). In both dimensions, the width selection is performed while ensuring that each region contains a sufficient number of data points. Specifically, the least number of data points within any region is 32, followed by 80, 84, etc. This

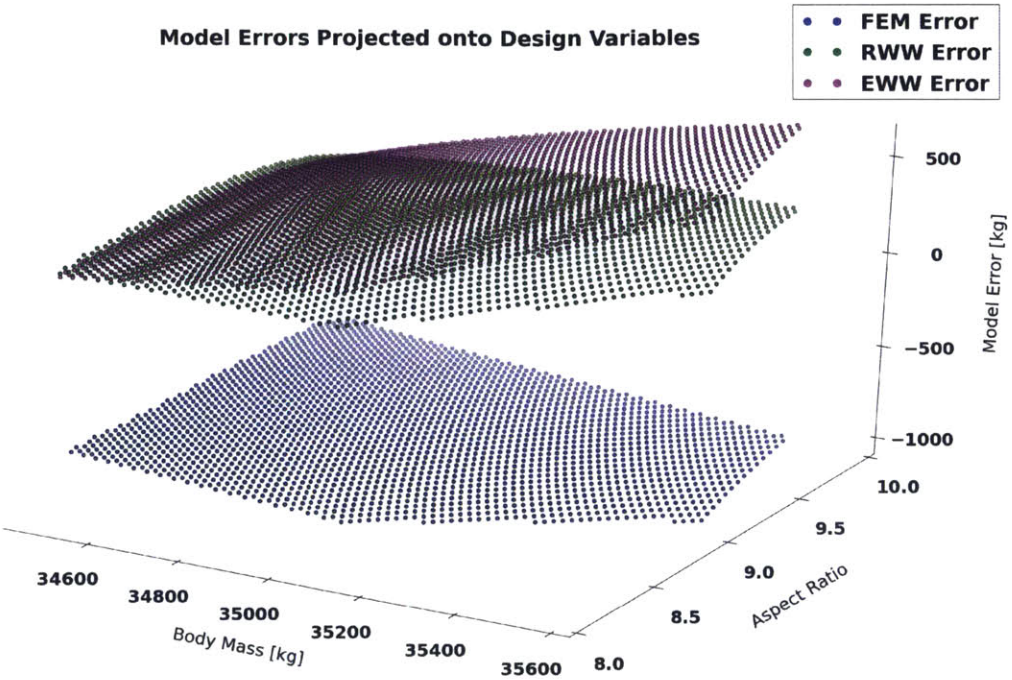

Figure 5-1: The colored points represent model discrepancies across the design space. Blue points correspond to the FEM, green points correspond to the RWW, and purple points correspond to the EWW. Points representing designs significantly different from the reference have been removed.

is enough for characterizing model correlation locally, and the result should be robust to outliers.

Regardless of how region widths are picked, it is essential that the range of design variables stay within the space where model correlation is defined. In the 2-D problem setup, the data for model correlation covers aspect ratios from 8.0 to 10.0 and body weights from 34500 kg to 36000 kg, while the corresponding design variables range from 8.40 to 9.60 and from 34890 kg to 35600 kg. Once the regional model correlations have been computed, bilinear interpolation is used to enforce continuity between adjacent regions. Figures 5-2, 5-3, and 5-4 show correlations between different pairs of models over both the maximum coverage of the data and the range of our design space.

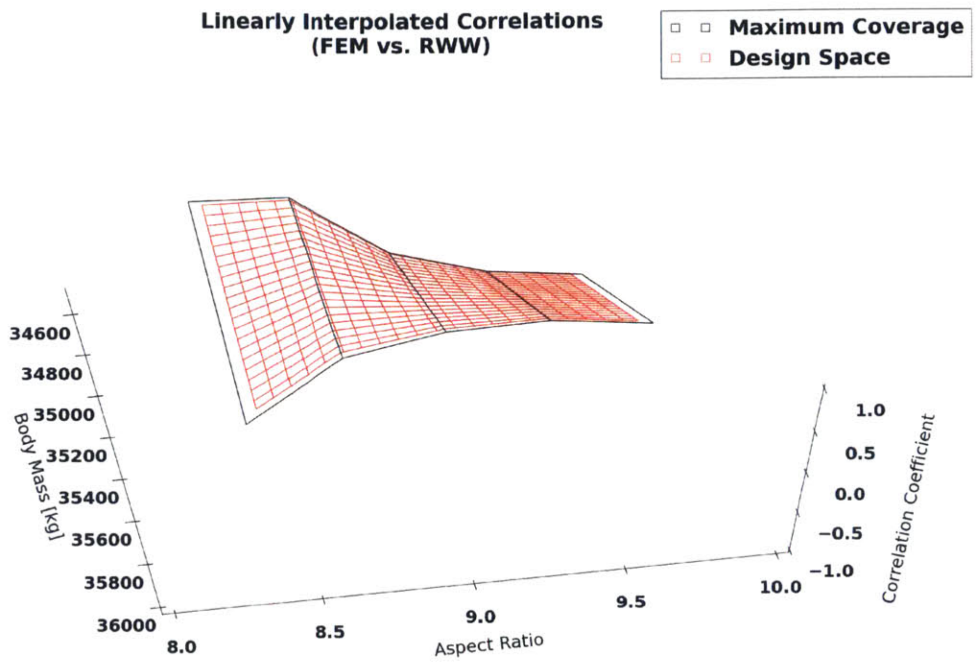

Figure 5-2: The piecewise linear surfaces show the interpolated model correlations between the FEM and the RWW. The grid outlined in black represents the maximum data coverage range, and the grid in red represents the range of the design space.

5.2.3 Relation to Model Physics

In order to validate the model correlations for information fusion, we must identify mechanisms driving the correlation trends. In 2-D, the complexity of our models and the multitude of error sources make it rather difficult to explain comprehensively all the relevant model physics and their interactions. Nevertheless, we provide some basic intuition for the observed correlation trends. In general, the greater the body weight, the higher the structural requirement for the wing (the aircraft requires more lift overall, which translates to higher moment at wing root). Under fixed wing span, varying aspect ratio only changes wing area, which translates to shifts in mean chord length. The FEM does not consider mean chord length, but for the RWW and the EWW, the shorter the chord, the higher the required wing spar cap thickness. In the arguments that follow, we will fix one dimension and discuss the correlation trend as the other dimension changes.

At low aspect ratio (8.0), the RWW's errors in wing weight prediction scales

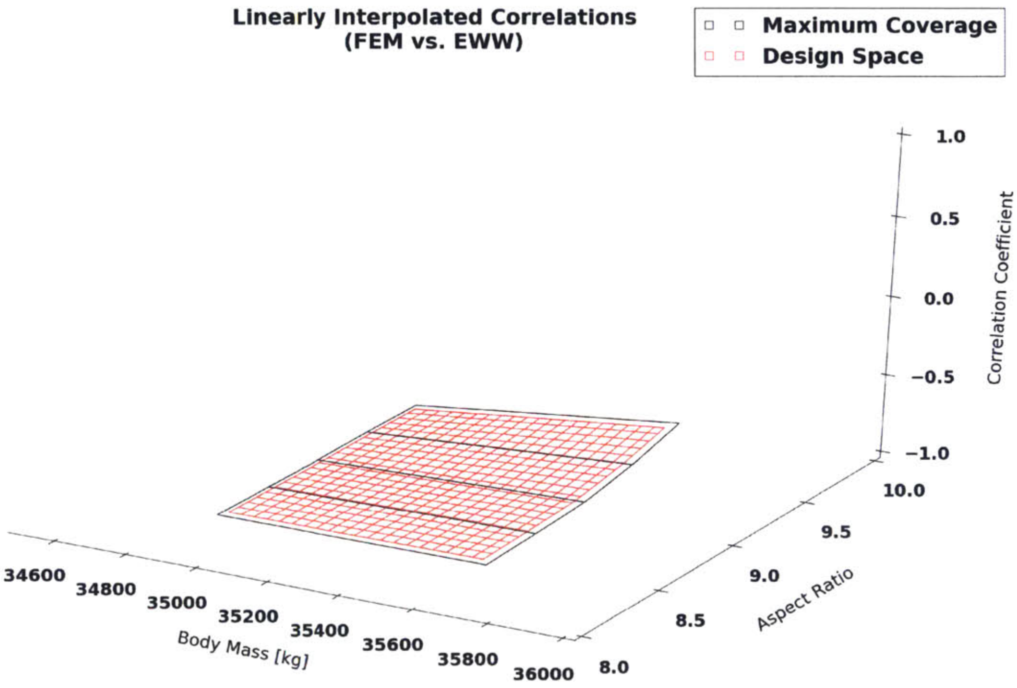

Figure 5-3: The piecewise linear surfaces show the interpolated model correlations between the FEM and the EWW. The grid outlined in black represents the maximum data coverage range, and the grid in red represents the range of the design space.

with the FEM's. This is because both the RWW and the FEM assume rectangular wing geometry, and at low aspect ratio, minor non-linearities in wing box component geometries present similar problems for both models (both models take weighted averages of dimensions drawn from fixed spanwise positions). However, these problems are more prominent in the FEM at high body weights, hence the fall in correlation as we move from low body weights to high body weights. At high aspect ratio (10.0), the RWW's errors in wing weight become negatively correlated with the FEM's. This is because higher aspect ratios are associated with shorter chord lengths, and a shorter chord means more spanwise non-linearities in wing box structural components. The FEM fails to capture these non-linearities to a large extent, but the RWW somewhat corrects for them through its internal sizing routine. At low body weight (34800 kg), the RWW and the FEM go from being positively correlated to being negatively correlated with increasing aspect ratio. This is because lower aspect ratios translate to longer chord lengths, which make the FEM's spanwise weighted averaging of structural components a more reasonable approximation. As aspect ratio increases,

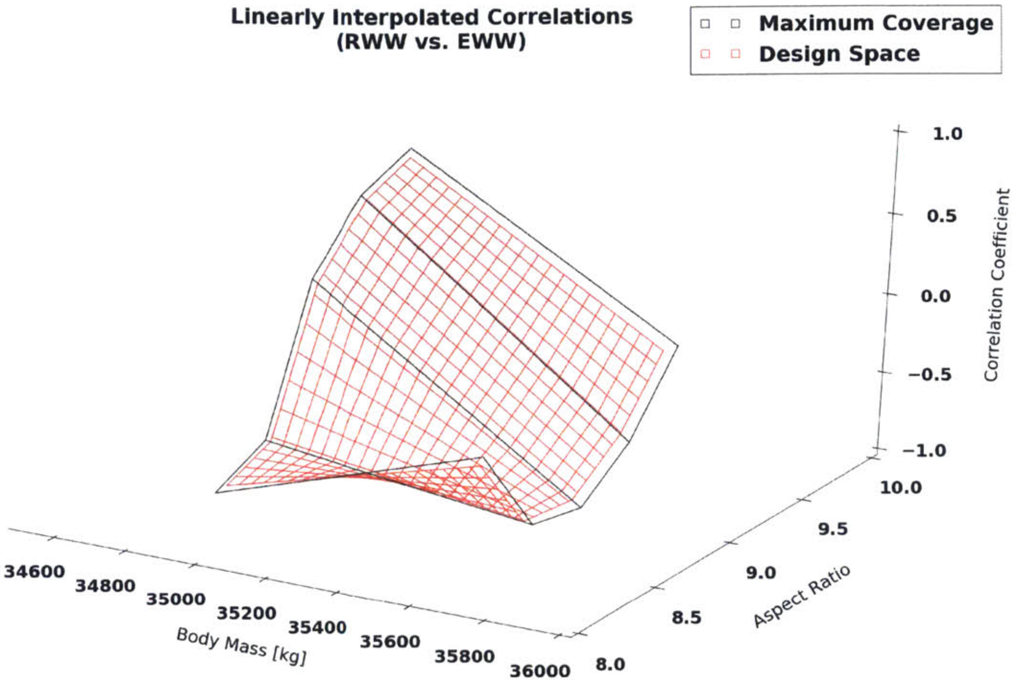

Figure 5-4: The piecewise linear surfaces show the interpolated model correlations between the RWW and the EWW. The grid outlined in black represents the maximum data coverage range, and the grid in red represents the range of the design space.

the decreasing chord length causes the FEM to incur worse errors from the resulting nonlinearities in geometry. In contrast, the RWW sizes the wing internally, thereby reducing the influence of the aforementioned nonlinearities. At high body weight (34800 kg), the RWW and the FEM stay highly positively correlated. This is because the RWW tends to underestimate the moment at wing root, and does so to a greater extent at higher body weights. The RWW also oversizes the wing spar caps, which, coupled with the moment underestimation, leads to underestimation in inboard weight and overestimation in outboard weight. Similarly, the FEM underestimates inboard weight and overestimates outboard weight by directly spanwise averaging the output wing geometry parameters from TASOPT.

At low body weight (34800 kg), the FEM and the EWW remain highly negatively correlated with increasing aspect ratio. This is because unlike the RWW, the EWW does not assume a fixed chord length throughout the wing (although it still assumes constant spar cap and web thicknesses). This somewhat corrects for the overestimation of outboard weight, but also leads to an overestimation of inboard weight.

Fuselage carryover further exacerbates this effect, so the EWW's error trends are the opposite that of the FEM's. At high body weight (34800 kg), the FEM and the EWW still remain highly negatively correlated with increasing aspect ratio. This is because at high body weight, the EWW's fuselage carryover correction leads to a more accurate wing root moment estimation, which translates to a structural overcompensation in the in-fuselage portion of the wing. The result is a large overestimation of inboard weight across all aspect ratios. In contrast, the FEM increasingly underestimates wing weight as it fails to account for nonlinearities in geometry.

At low body weight (34800 kg), the RWW and the EWW move from being negatively correlated to being positively correlated with as aspect ratio increases. This is because for the given body weight, the effect of over-designing the in-fuselage portion of the wing decreases with decreasing chord lengths (increasing aspect ratios). The RWW's errors trend downwards with increasing aspect ratio since its underestimation of wing root moment cancels out its over-sizing of spar caps as chord lengths become shorter. At high body weight (34800 kg), the RWW and the EWW errors stay negatively correlated with increasing aspect ratio, being most so near 8.5-9.0. This is because at high body weight, the errors derived from the EWW's overdesigning the in-fuselage portion of wing always fall towards overestimating wing weight regardless of chord lengths (shorter chord reduces weight in one direction, but causes structural requirements to go up in a perpendicular direction to compensate; the tradeoff becomes less favorable with increasing body weight). The RWW, on the other hand, always underestimates the wing root moment, which counteracts its overestimation of outboard weight by underestimating inboard weight. Its errors thus trend in the opposite direction to that of the EWW. At low aspect ratio (8.0), the RWW and the EWW stay negatively correlated with increasing body weight. This is because lower aspect ratio leads to longer mean chord, which worsens the effect of the EWW's oversizing the in-fuselage portion of the wing. In contrast, the RWW underestimates the wing weight consistently, leading to negative correlation between the two models. At high aspect ratio (10.0), the RWW and the EWW go from positively to negatively correlated with increasing body weight. This is because at low body weight, the extra moment exerted on wing root due to fuselage carryover does not contribute significantly towards the EWW's errors in oversizing the in-fuselage portion of the wing. Given that the RWW underestimates wing weight, both models have errors trending downwards. However, as body weight increases, the aforementioned extra moment becomes significant, so the EWW's errors vector towards overestimation while the RWW's errors vector towards underestimation.

5.3 Fidelity Variance Assignment

In the demonstration problem, fidelity variance is specified through the combination of a quadratic fit to model errors and a constant offset. In the 2-D problem, the quadratic fit is extended to both body weight and aspect ratio, yielding six coefficients. The coefficients correspond to the respective square terms, the interaction term, the respective linear terms, and the constant term. Figure 5-5 shows the quadratic fit to the respective model errors. The same constant offsets used in the demonstration problem are applied here, and the justifications for their selection translate as well. Specifically, we have 1500 kg for the FEM, 100 kg for the RWW, and 500 kg for the EWW. In Section 5.5, we will show that 1500 kg for the FEM is somewhat understated. We will also show the result from using 4000 kg instead of 1500 kg in order to illustrate the marginalization of the FEM's contribution to the fusion due to its severe inadequacy. Figure 5-6 shows the final model fidelity standard deviations after the constant offsets have been included.

5.4 Individual Model Surrogates

This section introduces the individual model surrogates produced using two different regression methods for the 2-D problem.

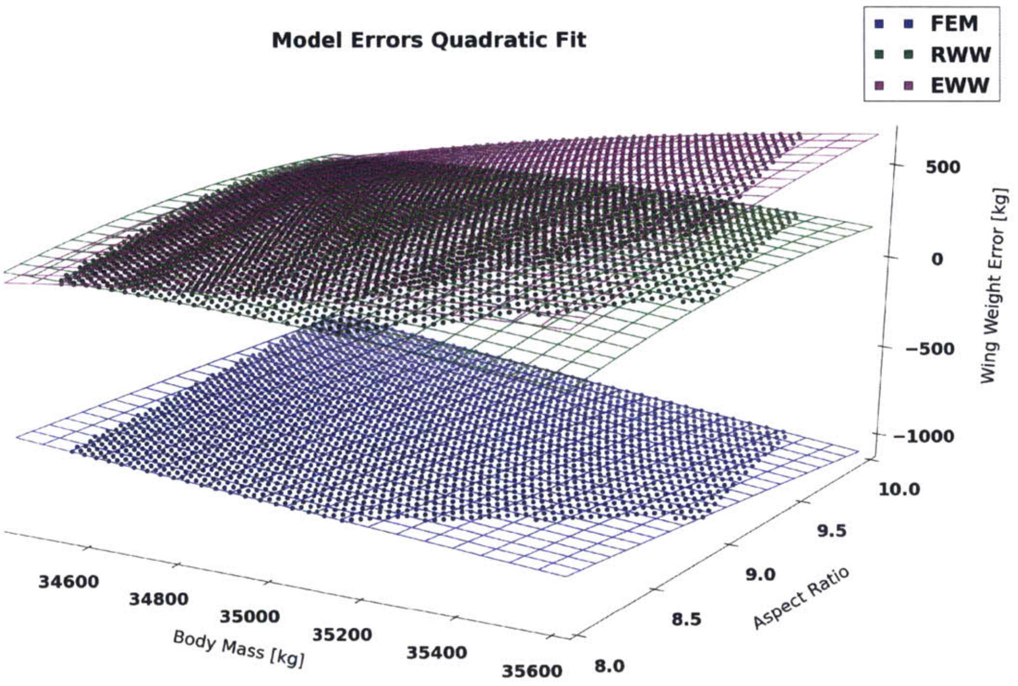

Figure 5-5: The colored points represent model discrepancies across the design space. The colored wireframes represent the quadratic fits to the respective model discrepancies. Blue color corresponds to the FEM, green color corresponds to the RWW, and purple color corresponds to the EWW.

5.4.1 Training Set

In the demonstration problem, we sampled each model sparsely and obtained training points that are more concentrated in certain areas. We generate the training set for the 2-D problem in a similar manner. The points are more densely packed where the models perform best, but at least one point lie in each extremity of the design space in order to guide the optimization of characteristic lengths. Note that in our particular 2-D problem, the design space is shaped as a rectangular plane. This is not true in general, since constraints may alter the feasible region of the original hyperplane.

For the FEM, the training points are picked at body weight [kg] and aspect ratio combinations of (35000, 8.67), (35000, 9.34), (35250, 9.0), (35500, 8.67), and (35500, 9.34). These five points are distributed symmetrically about the plane, such that one point sits at the center and one at each corner. The FEM considers neither body weight nor aspect ratio when computing wing weight, so its output is equivalent to the wing weight estimate of the reference aircraft (737-800) everywhere in the design

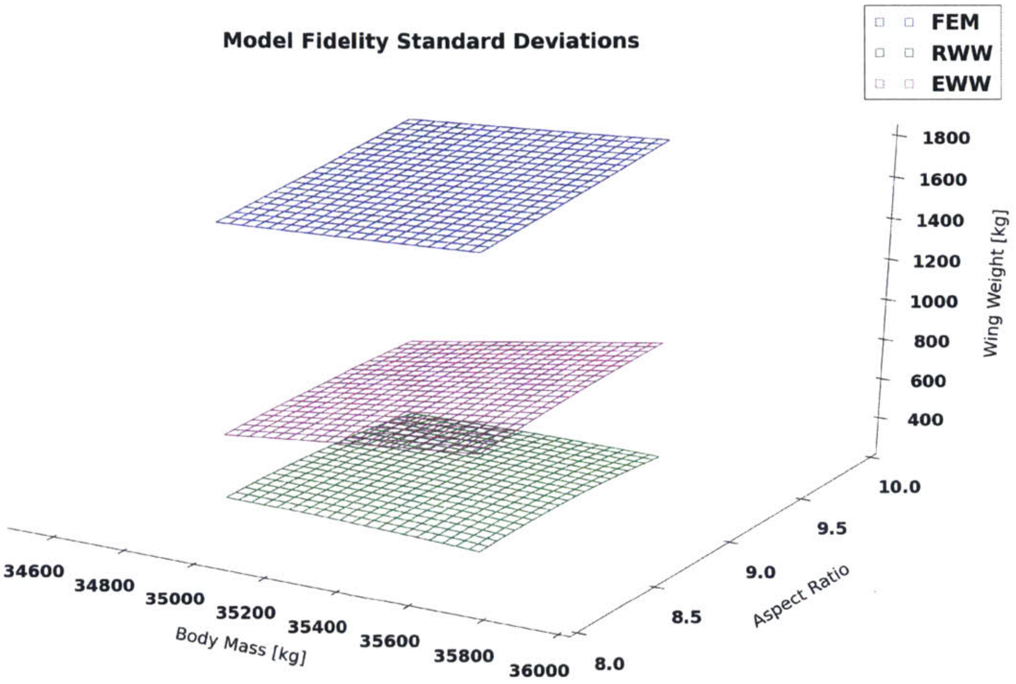

Figure 5-6: The colored wireframes represent the quadratic fits to model discrepancies after the addition of the constant offsets from expert elicitation. These values reflect uncertainty associated with model fidelity, and correspond to 1 Gaussian standard deviation.

space.

For the RWW, the training points are picked at body weight [kg] and aspect ratio combinations of (35000, 8.4), (35000, 9.0), (35500, 8.67), and (35500, 9.34). These four points are distributed around the plane with a bias towards low body weights and low aspect ratios. The RWW contains an internal sizing routine that uses both design variables when computing wing weight, and the direction where the sample concentration is biased indicate where the RWW's estimates are most accurate .

For the EWW, the training points are picked at body weight [kg] and aspect ratio combinations of (35000, 8.67), (35000, 9.34), (35000, 9.5), and (35500, 9.0). These four points are distributed around the plane with a bias towards low body weights and high aspect ratios. The EWW's fuselage carryover lift correction causes it to overestimate wing weight to a greater extent at higher body weights, so its samples are concentrated at low body weights. Given that the RWW already samples densely at low aspect ratios, we bias the EWW's training points towards high aspect ratios.

When devising a strategy for picking the training points, we also ensured that a response surface in the form of a plane can be fit at least deterministically via least squares. As such, both Gaussian Process Regression and Linear Regression can produce well-behaved surrogates.

5.4.2 Regression and Hyperparameters

Based on the training points from the previous section, we fit response surfaces using two different techniques. For the FEM, we fit a plane using Bayesian Linear Regression. This is an appropriate choice because the FEM's training points all share the same output value, and fitting a constant plane using Gaussian Process Regression will lead to numerical instabilities (characteristic lengths will be infinitely large). As mentioned in the demonstration problem, fusion of information can be conducted as long as the posterior distribution of the surrogate is Gaussian. Bayesian Linear Regression with an uninformative (zero mean, extremely large variance) prior conjugate to the Gaussian likelihood satisfies this requirement, as the Gaussian family is self-conjugate.

For the RWW, we fit a response surface using Gaussian Process Regression. The linearity in surface profile warrants Bayesian Linear Regression for generating the surrogate. However, in order to demonstrate fusion of information using surrogates produced by different techniques, we apply Gaussian Process Regression to the RWW's training set. As in the demonstration problem, we impose certain restrictions on hyperparameters. The sampling variance is bounded at 10000 kg, the characteristic length along body weight is bounded below at 100 kg, and the characteristic length along aspect ratio is bounded below at 0.1. Sampling variance is fixed to represent maximum uncertainty in the absence of nearby training points, where "nearby" corresponds to distances on the order of the respective characteristic lengths. The lower bounds on the characteristic lengths along body weight and aspect ratio are chosen to ensure that the model does not overfit the training set when picking hyperparameters through maximum likelihood estimate.

Similarly, we use Gaussian Process Regression for the EWW and apply the same restrictions on hyperparameters. Figures 5-7, 5-8, and 5-9 show the respective means of the individual surrogates in conjunction with the training set that produced them. Figure 5-10 presents a side-by-side comparison of the means, and Figure 5-11 with the standard deviations. Note that the standard deviations from Figure 5-11 include sampling uncertainty, so they are greater than those given in Figure 5-6.

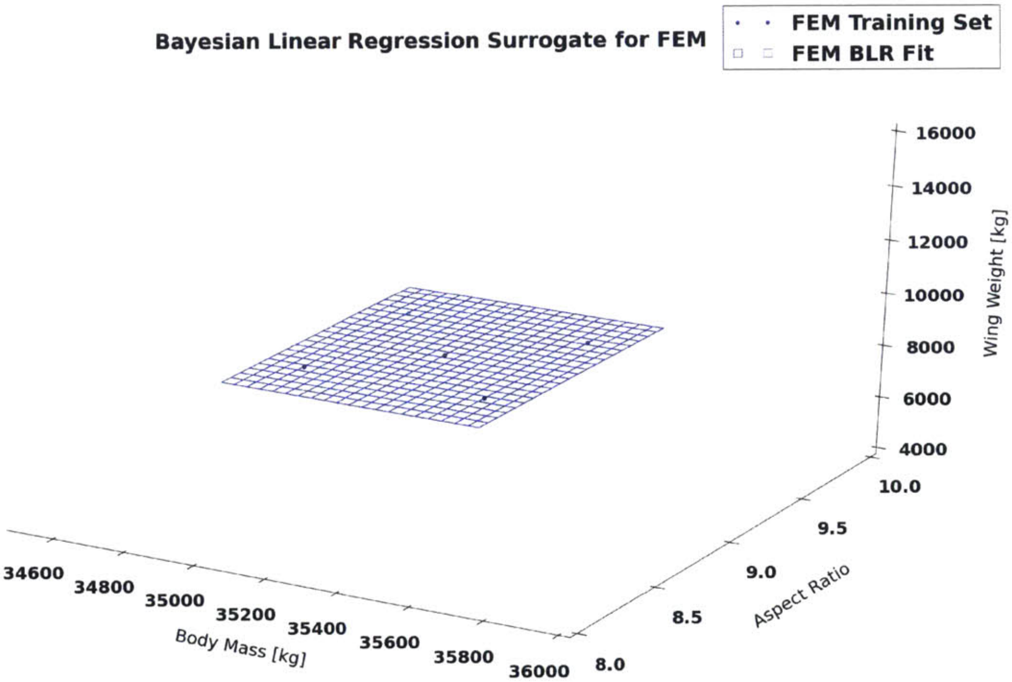

Figure 5-7: The blue points correspond to the training points for the FEM's surrogate. The blue wireframe represent the FEM's surrogate mean. The FEM's surrogate is produced using Bayesian Linear Regression.

Optimization of hyperparameters by maximum likelihood yields the following results: where the characteristic lengths along both dimensions for the RWW and the

	Sampling Variance σ_{GP}^2 $[kg^2]$	Characteristic Body Weight l_{bw} $[m]$	Characteristic Aspect Ratio l_{AR}
The FEM	N/A	N/A	N/A
The RWW	10000	56251	5.74
The EWW	10000	61300	6.39

Table 5.1: The columns show the respective Gaussian Process Regression hyperparameters for the 2-D problem.

EWW suggest that the surface profiles are linear over the design space. The surro-

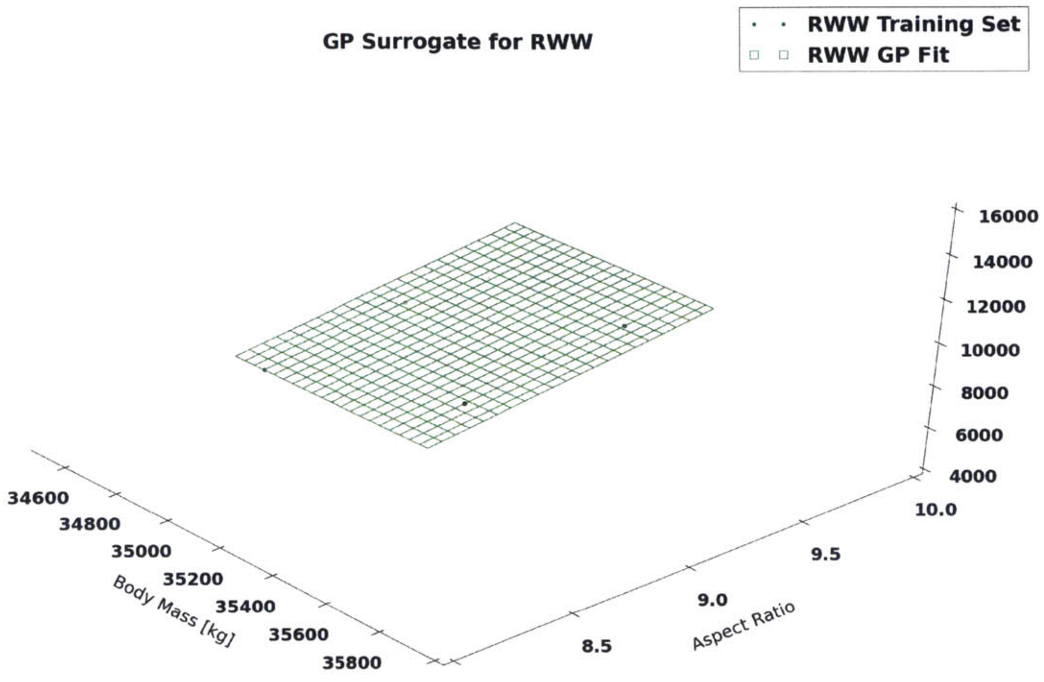

Figure 5-8: The green points correspond to the training points for the RWW's surrogate. The green wireframe represent the RWW's surrogate mean. The RWW's surrogate is produced using Gaussian Process Regression.

gate for the FEM has no hyperparameters because it is not produced using a kernel method. The coefficients of the linear regression are 0 for each of the the linear terms and 9006 for the constant term. The model means follow similar trends, with the RWW and the EWW almost overlapping and the FEM a short distance below. With regard to standard deviations, the FEM has the highest out of all three, the EWW the second highest, followed by the RWW. In this particular scenario, there is a clear hierarchy in model fidelity, as the RWW outperforms the EWW, and the EWW outperforms the FEM everywhere in the design space. The formation of the hierarchy is due to expert elicitation, and does not happen in general. Furthermore, Figure 5-11 and Figure 5-6 appear almost identical. This is because uncertainty from sampling is almost negligible compared to uncertainty from model fidelity. The linearity of the surface profile greatly diminishes the effect of σ^2_{GP} on the overall model uncertainty, so model fidelity becomes the dominant factor in deciding the reliability of the surrogate.

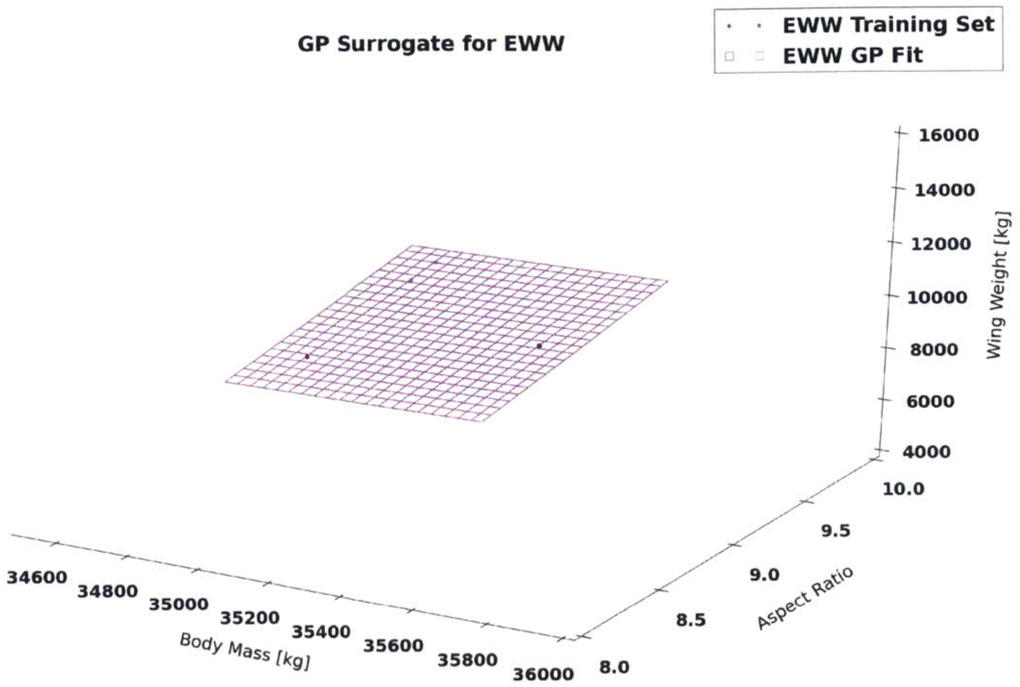

Figure 5-9: The purple points correspond to the training points for the EWW's surrogate. The purple wireframe represent the EWW's surrogate mean. The EWW's surrogate is produced using Gaussian Process Regression.

5.5 Information Fusion

This section reviews the results of information fusion with model correlation.

5.5.1 Results

We present the results for information fusion under model independence in the same plot as information fusion with model correlation to provide a basis of comparison. Figure 5-12 shows the respective fused means, and Figure 5-13 shows the respective standard deviations.

The fused mean for the correlated case mostly lies below that for the independent case. The slope of the surface is also different for the correlated case compared to the independent case. This is because the contribution of the FEM to the fused surrogate is overstated. Since the FEM's surrogate mean lies below that of the other two models and its surface harbors no slope, the information it carries conflicts with those supplied by its counterparts. Given that the FEM is the most inadequate model out of the

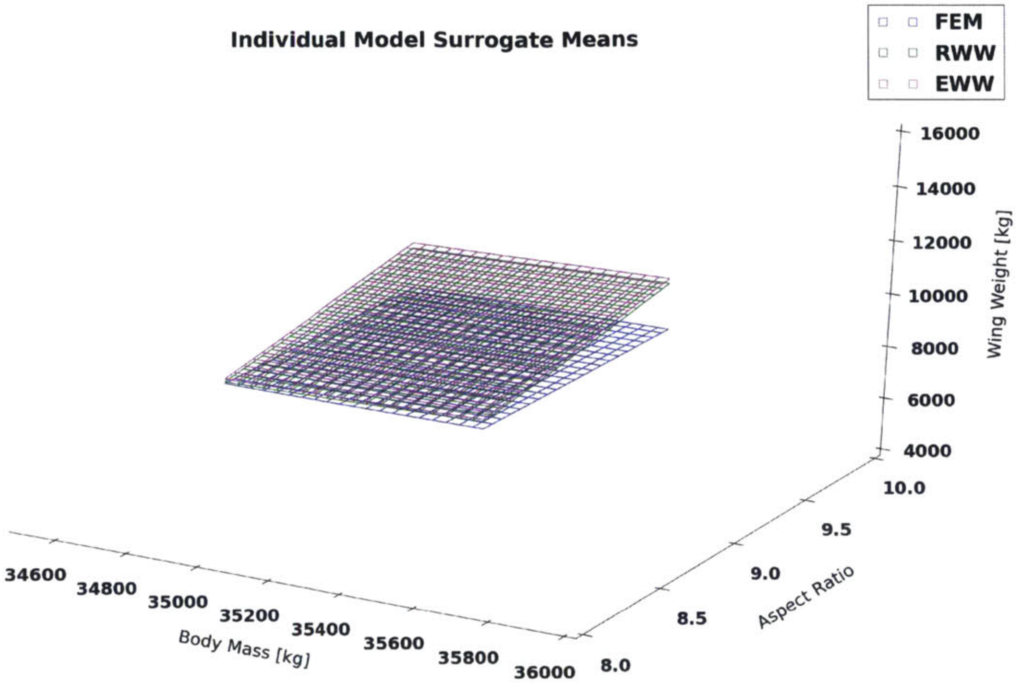

Figure 5-10: The blue wireframe represents the FEM's surrogate mean, the green wireframe represents the RWW's surrogate mean, and the purple wireframe represents the EWW's surrogate mean.

three (it does not use either of our design variables in computing wing weight), it is likely that the level of fidelity it was assigned during expert elicitation had been overstated. Coupled with the influence from model correlation, the contribution from the FEM causes the fused surrogate to be behave inconsistently. Specifically, there is an unjustified sharp increase in the fused mean towards high body weight and aspect ratio.

The fused standard deviation for the correlated case is lower than that for the independent case. The cause for the improved overall confidence lies in the negative correlation between the FEM and the EWW. Taking advantage of the opposing error trends between the FEM and the EWW, the fused surrogate can more readily infer the direction of the true output. The fused standard deviation troughs at low body mass and high aspect ratio. This is because both the RWW and the EWW have lowest uncertainty there, and the models are mostly independent in that region. The fused standard deviation peaks near medium body mass and high aspect ratio. This is due to complex interactions among model correlations and uncertainties, such that

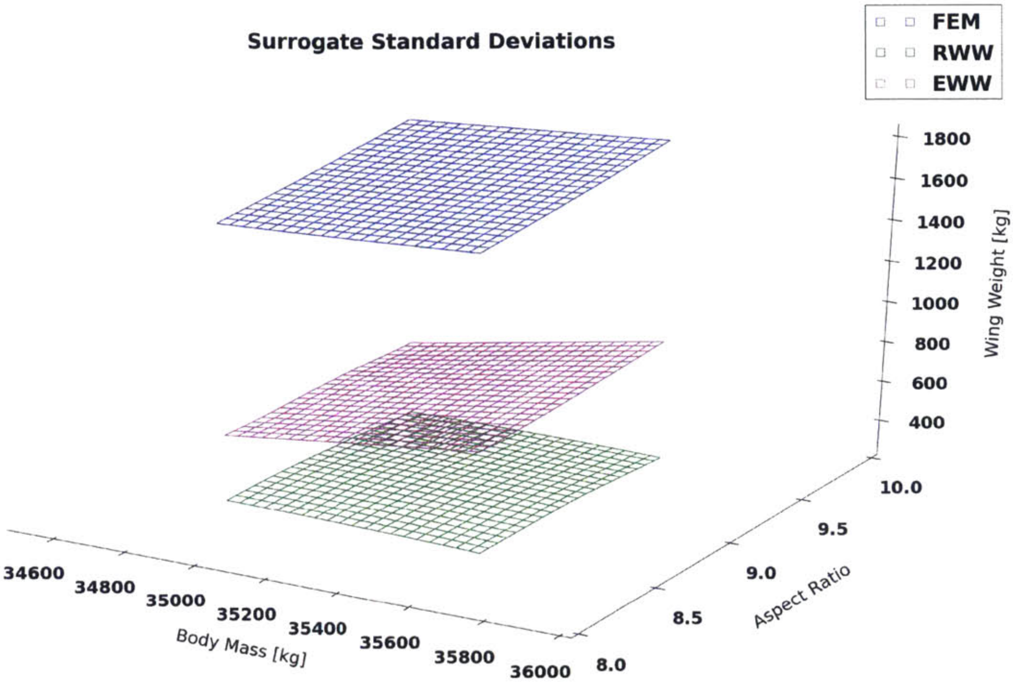

Figure 5-11: The blue wireframe represents the FEM's surrogate standard deviation, the green wireframe represents the RWW's surrogate standard deviation, and the purple wireframe represents the EWW's surrogate standard deviation. The standard deviations include both uncertainty from sampling and uncertainty from model fidelity.

the variance inflation due to inconsistency of information outweights the variance reduction from extrapolation based on model correlation.

To improve the representation of the FEM's model fidelity, we add 4000 kg of constant offset instead of the original 1500 kg to the quadratic fit of the FEM's errors. The new information fusion results are shown in Figures 5-14 and 5-15.

Under the FEM's updated model fidelity, the fused surrogate becomes much more sensible. The fused means of the two cases are more similar, aside from the slightly increased variability in surface profile for the correlated case. The contribution from the FEM has been marginalized; the fused surrogate contains less information from the FEM. However, at high body weight and high aspect ratio, the fused mean is still pushed upwards. This can be attributed to the negative weight associated with the FEM's mean in the fusion formula. The weight is negative because the FEM and the RWW are highly correlated there, and the FEM carries the highest variance out

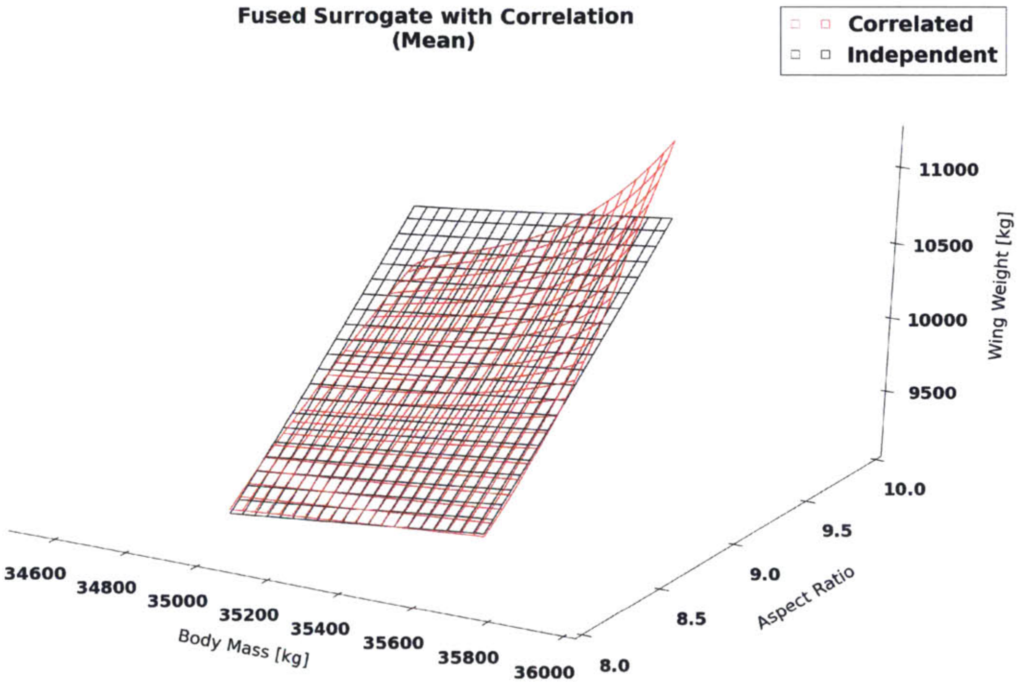

Figure 5-12: The red wireframe corresponds to the mean of the fused surrogate under model correlation. The black wireframe corresponds to the mean of the fused surrogate under model independence.

of the three models (the difference in variance between the FEM and the RWW is greater than the difference in variance between the FEM and the EWW, so the FEM's highly positive correlation with the RWW overpowers its highly negative correlation with the EWW in dictating its weight in the fused mean, making it negative overall). Nevertheless, we have shown that by adding uncertainty to a model, we can reduce its contribution to the fusion of information. In the most extreme case, one can simply remove the model from consideration altogether, which is perhaps more sensible (and computationally beneficial) than adding a large uncertainty to its output. The fused standard deviation also exhibits a small overall increase, since the FEM can no longer be relied upon as an accurate information provider.

5.5.2 Analysis and Discussion

The fused standard deviation under model correlation forms arches near each edge of the design space and exhibits significant variation. In contrast, the fused standard

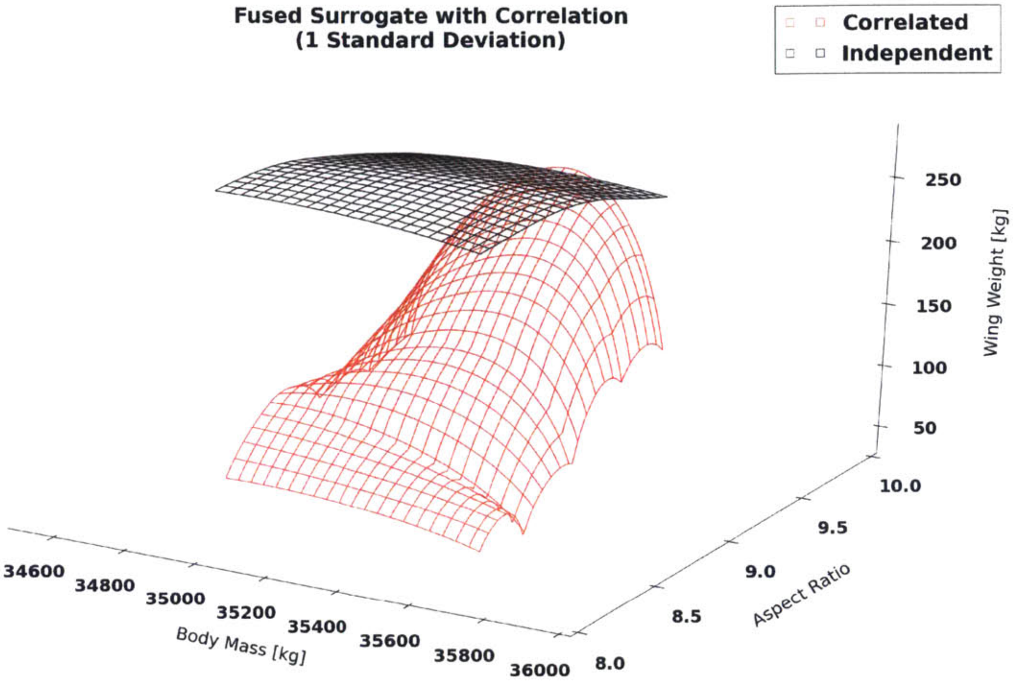

Figure 5-13: The red wireframe corresponds to the standard deviation of the fused surrogate under model correlation. The black wireframe corresponds to the standard deviation of the fused surrogate under model independence.

deviation under model independence appears generally flat. For the fused standard deviation under model correlation, the largest arch occurs along aspect ratio of 9.60, and the top of the arch lies near the body weight of 35200 kg. On first glance, since Gaussian Process Regression often produces surrogates that have arch-shaped sampling standard deviations with roots at training points, those features may readily propagate into the standard deviation of the fused surrogate. While the sampling uncertainties do play a role, their contribution to the fused standard deviation turns out to be negligibly small. The primary contribution comes from uncertainties associated with model fidelity. In addition, model correlations dictate the reduction or inflation of standard deviations during information fusion. Consequently, the roots of the arches in the fused standard deviation under model correlation do not coincide with positions of any of the constituent surrogates' training points, but instead coincide with edges in the surfaces representing model correlations.

Recall the information fusion formulae given by Equation 4.2, and recall that the fused variance at any given point is computed by summing up all elements in Σ^{-1}

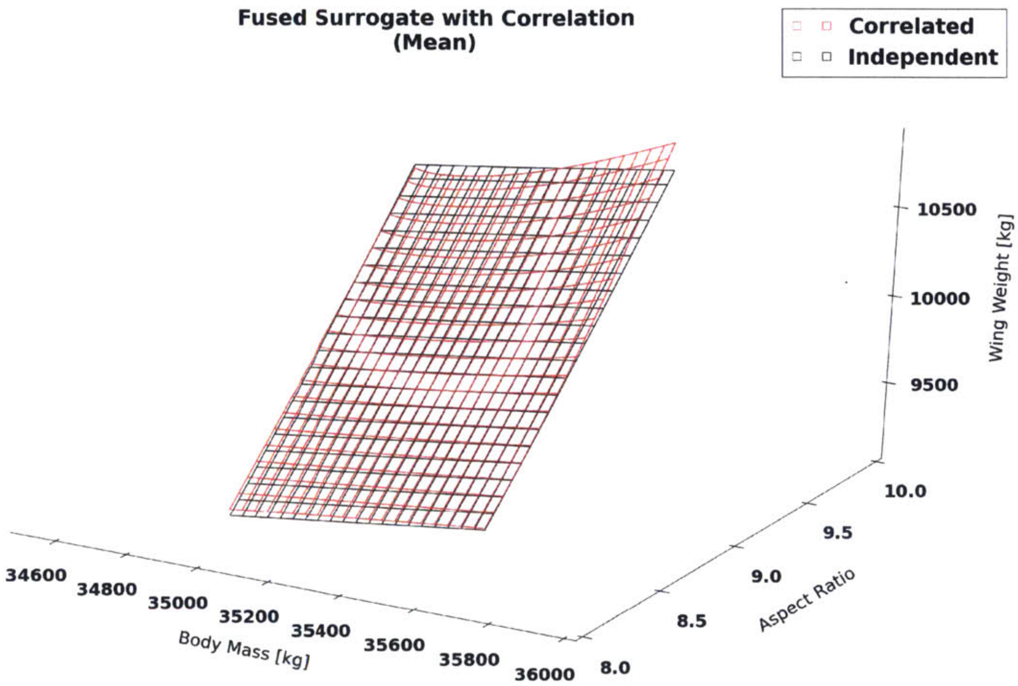

Figure 5-14: The red wireframe corresponds to the mean of the fused surrogate under model correlation, after the FEM's fidelity has been updated. The black wireframe corresponds to the mean of the fused surrogate under model independence, after the FEM's fidelity has been updated.

and taking the reciprocal of the result. The larger the sum of all elements in Σ^{-1} (corresponding to the denominators of the mean and variance terms in Equation 4.2), the smaller the fused variance. Figure 5-16 shows the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ (the sum of all elements in Σ^{-1}) across the design space. Note that all of our subsequent analyses are conducted based on the updated model fidelity for the FEM (4000 kg in constant offset).

As expected, the peaks of Figure 5-16 matches the troughs of Figure 5-15. But why would $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ exhibit any peaks in the first place? And why would $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ not vary linearly, given that linear interpolation is used to enforce continuity of model correlation between adjacent regions? Furthermore, why would certain peaks occur at positions where the correlation between none of the pairs of the models appear to be at an extreme (possess values close to -1 or 1)?

Fortunately, all of these questions can be answered by inspecting Figure 5-17. Figure 5-17 shows corresponding values of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ over all combinations of model

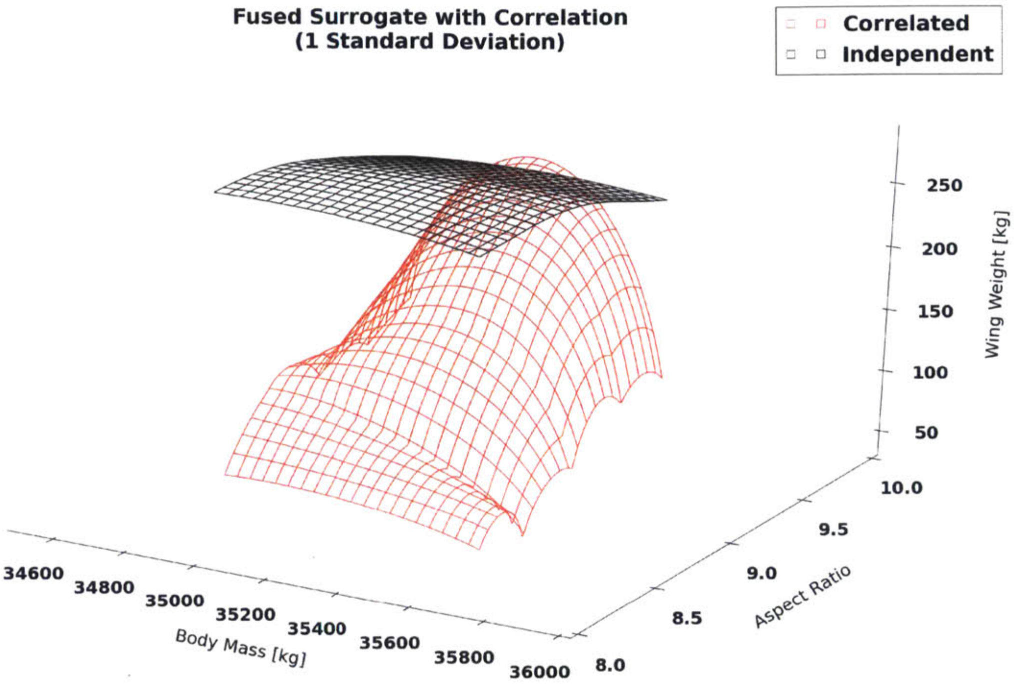

Figure 5-15: The red wireframe corresponds to the standard deviation of the fused surrogate under model correlation, after the FEM's fidelity has been updated. The black wireframe corresponds to the standard deviation of the fused surrogate under model independence, after the FEM's fidelity has been updated.

correlations. Figures 5-18, 5-19, and 5-20 show slices along the x, y, and z axes of the 3-D geometry in Figure 5-17 respectively. In the volumetric plot, color represents the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$—the brighter the color, the larger the value. The x, y, and z axes (the z-axis is vertical) correspond to the correlations between the FEM and the RWW, the FEM and the EWW, and the RWW and the EWW, respectively. Furthermore, we have restricted the plot to show only results for positive-definite Σ matrices. That is, any combinations of model correlations yielding non positive-definite Σ matrices are ignored, since they lead to negative fused variances that carry no physical significance. An interesting corollary is that the volume enclosed by the colored surface in Figure 5-17 also represents the range of all positive-definite 3 by 3 Σ matrices (under fidelity standard deviation constant offsets of 4000 kg for the FEM, 100 kg for the RWW, and 500 kg for the EWW). In general, the closer the correlation between a given pair of models is to 1 or -1, the less room for variation the other model correlations have in order to maintain positive-definiteness in Σ.

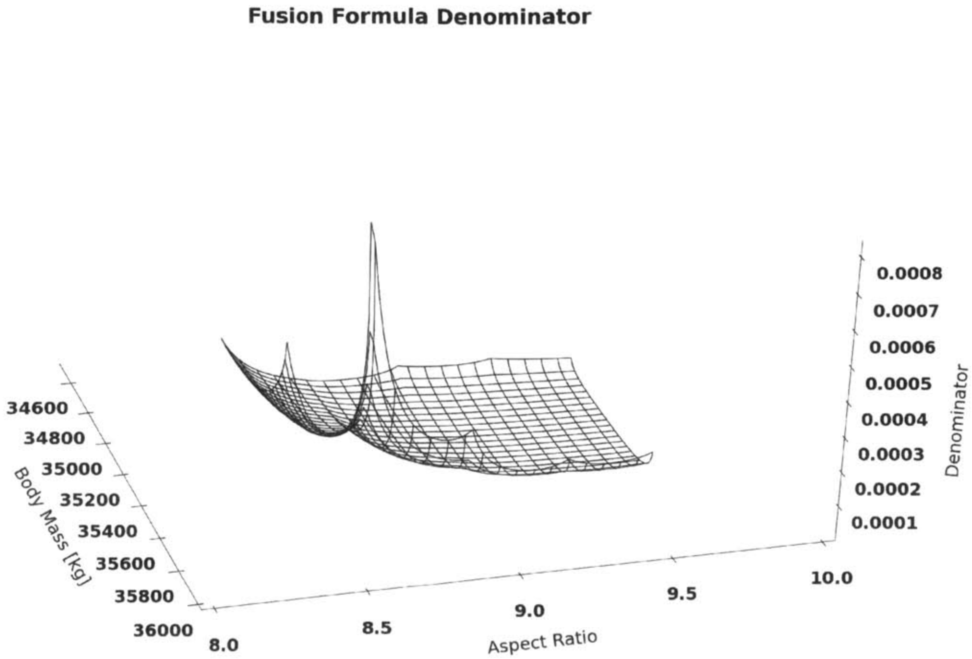

Figure 5-16: The black wireframe represents the quantity $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ (corresponding to the denominators of the mean and variance terms in Equation 4.2) across the design space. The z-axis uses linear scale.

Peaks in $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ occur at certain combinations of model correlations. In particular, these peaks are almost exclusive to combinations involving negative correlations. In the 2-D problem, the FEM and the EWW are highly negatively correlated everywhere. Taking slices along the y-axis (the correlation between the FEM and the EWW) as showing in Figure 5-19, we see that peaks tend to form as correlation between the FEM and the RWW approaches -1, or as correlation between the RWW and the EWW approaches -1, or through some combination of both. A correlation close to -1 between one or more pairs of models will overpower moderate correlation between the remaining pairs. For instance, if the correlation between the FEM and the EWW and the correlation between the RWW and the EWW are both close to -1, then a peak will form even if the correlation between the FEM and the RWW is forced to take on a moderately positive value in order to maintain positive-definiteness of the overall correlation matrix. In fact, this example explains the highest peak in Figure 5-16. However, a correlation close to -1 between a pair of models is not required for a peak to form, as combinations of non-extreme correlation values can still specify

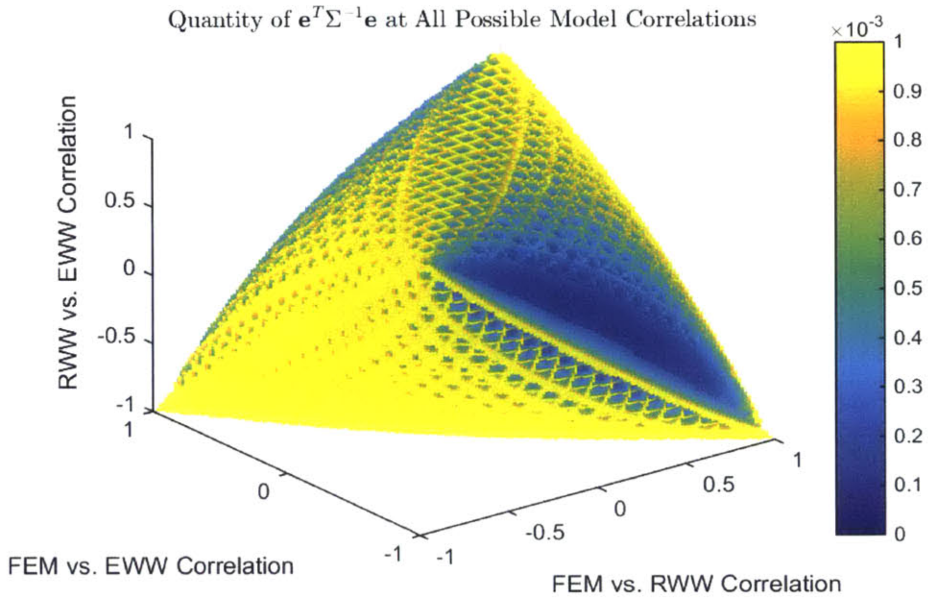

Figure 5-17: The 3-D geometry encompasses the set of all positive-definite 3 by 3 Σ matrices (under fidelity std. offsets of 4000 kg for the FEM, 100 kg for the RWW, and 500 kg for the EWW). The colors represent the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ corresponding to different combinations of model correlations. Yellow signifies values close to 0.0010, blue signifies values close to 0.

points that lie in the peak regions (denoted in Figure 5-17 by bright yellow color).

Figure 5-17 also reveals the reason behind $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ not varying linearly. Given any combination of model correlations, we can identify a value based on color from Figure 5-17, corresponding to the associated $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ quantity. A linear interpolation between correlation matrices $\boldsymbol{\rho}_1$ and $\boldsymbol{\rho}_2$ corresponds a linear translation in Figure 5-17 between points with coordinates specified by model correlation values in $\boldsymbol{\rho}_1$ and $\boldsymbol{\rho}_2$. Points that lie on the interpolated segment correspond to points that lie along the linear translation path. Judging from Figure 5-17, changes in $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ along any linear translation path do not have to be linear. Depending on the direction of translation, $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ may nonlinearly increase, decrease, or even reverse direction somewhere in the middle (thanks to local convexity of the domain). Figure 5-16 shows the complex behavior of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ specific to the 2-D problem.

Aside from model correlation, the respective model uncertainties also contribute to the formation of peaks. Figure 5-21 shows corresponding values of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ for all

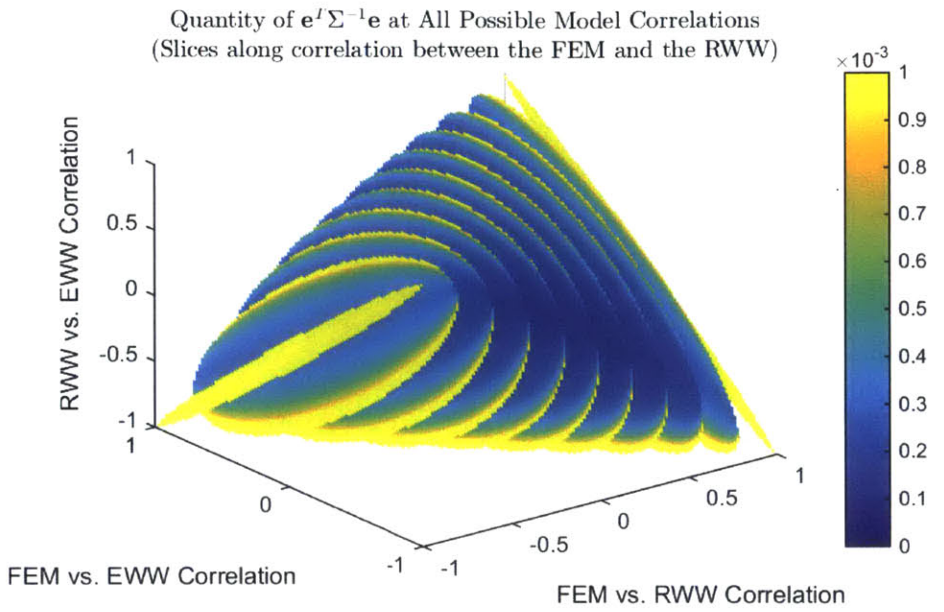

Figure 5-18: The layers represent slices along the x-axis (corresponding to the correlation between the FEM and the RWW) of the 3-D geometry in Figure 5-17. The colors represent the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ corresponding to different combinations of model correlations. Yellow signifies values close to 0.0010, blue signifies values close to 0.

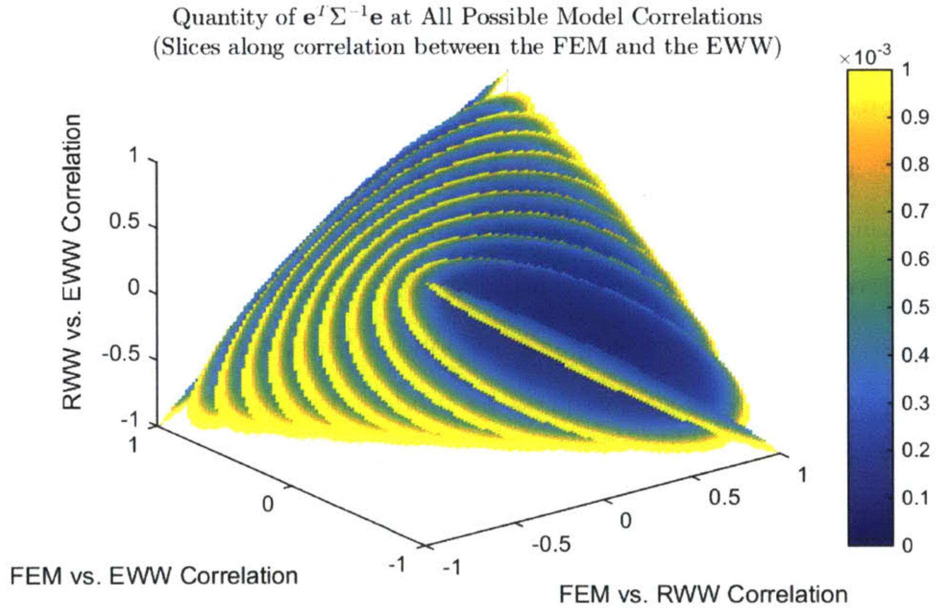

Figure 5-19: The layers represent slices along the y-axis (corresponding to the correlation between the FEM and the EWW) of the 3-D geometry in Figure 5-17. The colors represent the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ corresponding to different combinations of model correlations. Yellow signifies values close to 0.0010, blue signifies values close to 0.

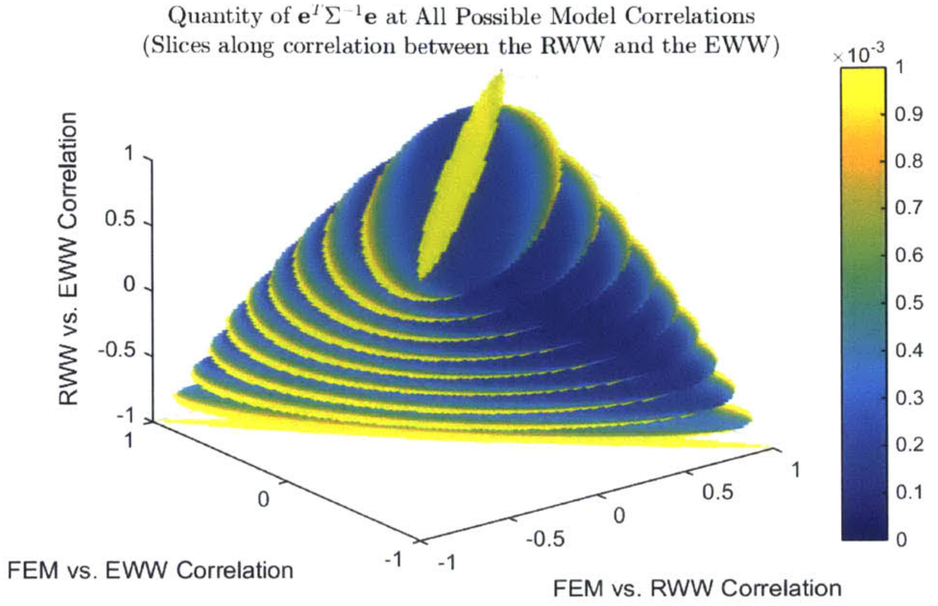

Figure 5-20: The layers represent slices along the z-axis (corresponding to the correlation between the RWW and the EWW) of the 3-D geometry in Figure 5-17. The colors represent the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ corresponding to different combinations of model correlations. Yellow signifies values close to 0.0010, blue signifies values close to 0.

combinations of model correlations when all the models have equal fidelity standard deviations of 100 kg. Figures 5-22, 5-23, and 5-24 show slices along the x, y, and z axes of the 3-D geometry in Figure 5-21 respectively.

Since large $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ values correspond to small fused standard deviations, a negative correlation between two models leads to the extraction of more information than either model independently carries when performing information fusion. This extrapolation can be interpreted as an inference based on model discrepancy trends, which provides additional information on top of the contribution from each constituent model. Nevertheless, confidence in the fused result is contingent upon thorough understanding of all the model error trends, the responsible model physics, and the proper reasoning behind all the model correlations. As such, certain combinations of model correlations can vastly reduce uncertainty in the final surrogate. However, we must caution that the fused standard deviations are likely understated, since we have not accounted for uncertainty in the model correlations themselves.

As in the demonstration problem, we provide the condition numbers of both Σ and

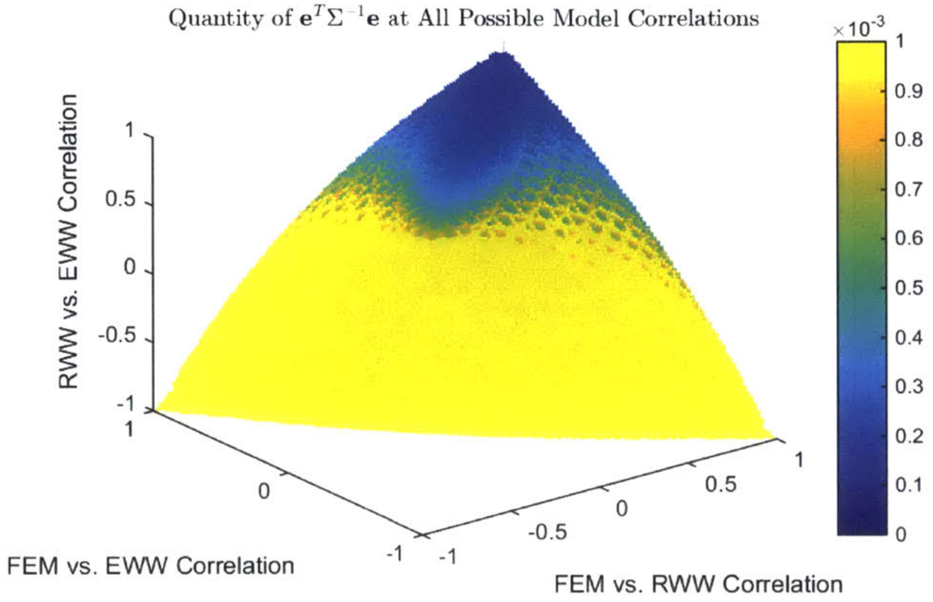

Figure 5-21: The 3-D geometry encompasses the set of all positive-definite 3 by 3 Σ matrices (under fidelity std. offsets of 100 kg for the FEM, 100 kg for the RWW, and 100 kg for the EWW). The colors represent the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ corresponding to different combinations of model correlations. Yellow signifies values close to 0.0010, blue signifies values close to 0.

$\boldsymbol{\rho}$ matrices for verification. Figure 5-25 shows the result. Visual inspection reveals that the maximum condition number lies around 10^4, which is sufficiently low for our purposes.

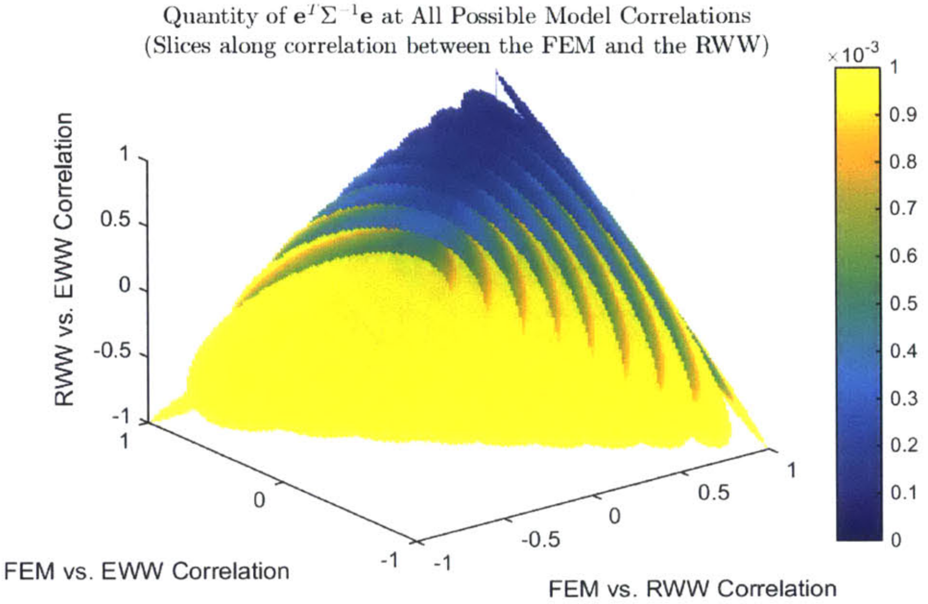

Figure 5-22: The layers represent slices along the x-axis (corresponding to the correlation between the FEM and the RWW) of the 3-D geometry in Figure 5-21. The colors represent the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ corresponding to different combinations of model correlations. Yellow signifies values close to 0.0010, blue signifies values close to 0.

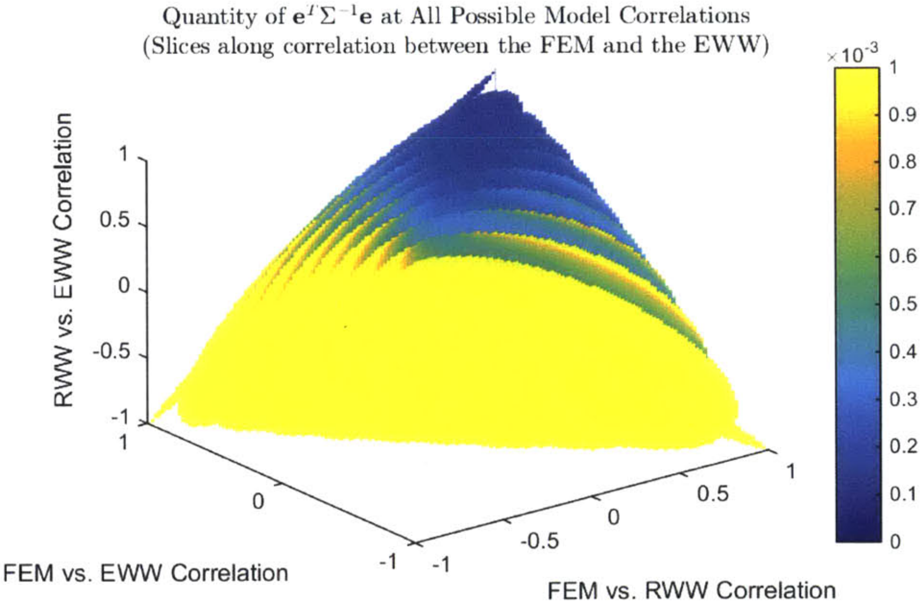

Figure 5-23: The layers represent slices along the x-axis (corresponding to the correlation between the FEM and the EWW) of the 3-D geometry in Figure 5-21. The colors represent the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ corresponding to different combinations of model correlations. Yellow signifies values close to 0.0010, blue signifies values close to 0.

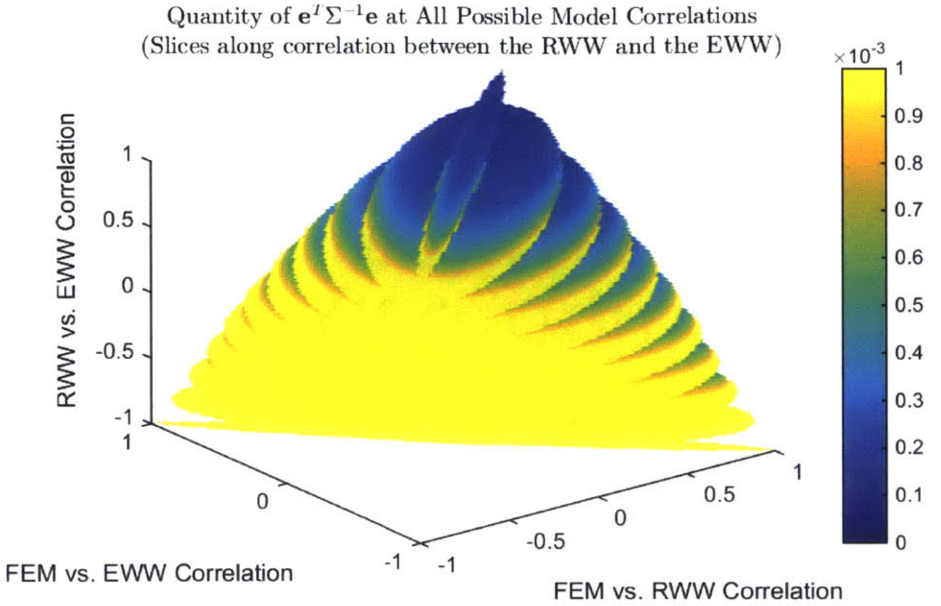

Figure 5-24: The layers represent slices along the x-axis (corresponding to the correlation between the RWW and the EWW) of the 3-D geometry in Figure 5-21. The colors represent the value of $\mathbf{e}^T\Sigma^{-1}\mathbf{e}$ corresponding to different combinations of model correlations. Yellow signifies values close to 0.0010, blue signifies values close to 0.

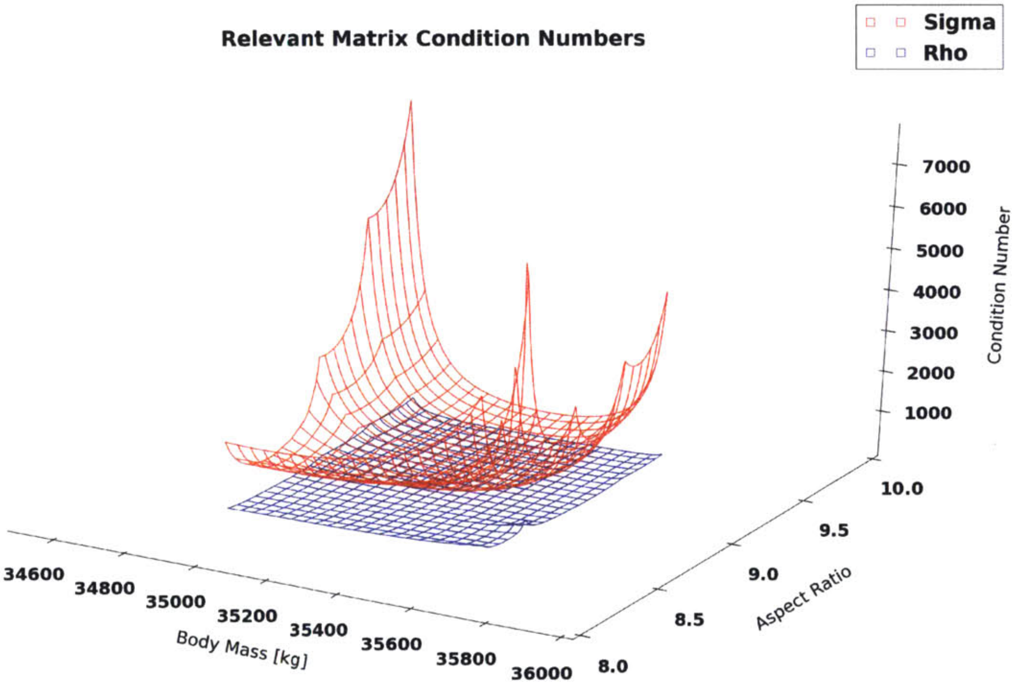

Figure 5-25: The red wireframe represents the condition number of Σ across the design space. The blue wireframe represents the condition number of ρ across the design space. Notice that the y-axis uses linear scale.

Chapter 6

Conclusion

This chapter concludes our development of the multifidelity framework. Section 6.1 highlights the main features of the proposed framework. Section 6.2 supplies directions with potential for future work.

6.1 Summary

The primary objective of this work is to incorporate model correlation when combining information from different sources. In aid of the framework's development, we created three models of varying fidelities, the Finite Element Model (FEM), the Rectangular Wing Weight (RWW) model, and the Elliptic Wing Weight (EWW) model. The fidelity levels of the three models change as functions of position in the design space. Using TASOPT as truth model and 737-800 as reference aircraft, we set up a 1-D demonstration problem with wing span as design variable and wing weight as predicted quantity of interest.

Model discrepancy data over the range of the design space is required for quantifying model correlations and fidelities. To produce the data, we perturb the parameters of the reference aircraft and use them as inputs to each model. Model discrepancies are then computed by taking differences between the output of each model and the output of TASOPT. In order to track model correlation across the design space, we divide it into different regions each containing some minimum number of data points.

Region widths are chosen based on model physics, and any variability in model correlation is connected to some underlying mechanism behind the models. Once regional model correlations are defined, we apply linear interpolation between adjacent region centroids to enforce both continuity and positive-definiteness in correlation matrices.

The data used for characterizing model correlation is also used to estimate model fidelity. Expert elicitation indicates that quadratic fits to model errors followed by various degrees of constant offset are sufficient for approximating associated output variances. Once model fidelities have been assessed, we designate positions in the design space where training points for respective model surrogates are to be evaluated. Bayesian Regression techniques are used to build surrogates that contain uncertainty from sampling. Summing together uncertainty from both model fidelity and sampling yields overall surrogate variance.

Fusion of information is conducted both with and without consideration for model correlation. The weight vector pre-multiplying the component surrogate means serves as a basis for weighing the contribution of each constituent surrogate towards the fused surrogate. In the demonstration problem, fusion under model correlation produces a different result from fusion under model independence. Two conflicting forces govern the behavior of the fused variance. The first is variance reduction from extrapolated information based on model correlation, and the second is variance inflation due to inconsistency of information carried by highly correlated models.

The same procedure is applied to a problem that considers two design variables of body weight and aspect ratio. The results reveal deeper insight into the behavior of the fused variance in relation to different combinations of pairwise model correlations. In general, negative model correlations allow for the extraction of more information from each model surrogate than it independently carries. This extrapolation can be interpreted as an inference based on model correlations, which provides additional information on top of the contribution from each model alone. Hence, there is potential in boosting the confidence of the fused surrogate when model correlation is accounted for.

6.2 Future Work

The proposed multifidelity framework opens up a few potential areas for further exploration. Topics mentioned below are aimed at extending the functionality and improving the performance of the framework.

6.2.1 Systematic Region Width Selection

In the presented framework, we stratify data for characterizing model correlation according to manually selected region widths. The region widths are justified through knowledge of model physics. However, in reality, we often do not have as detailed information regarding a model's underlying mechanisms as we do in the demonstration problem. Therefore, it would expand the applicability of our framework if there can be a systematic method for determining region widths. A starting recommendation is perhaps to consider the sensitivity of each design variable to the model output and choose region widths such that the same percentage change in model output can be produced by traversing the same units of region width in any dimension. However, this method depends on the monotonicity of the model output along all dimensions, and further study is needed before it can be integrated into the framework.

6.2.2 Similarity Functions

Data for characterizing model correlation in the 2-D problem is trimmed via expert elicitation based on similarity to the reference aircraft, where similarity is measured by Euclidean distance across a select number of parameters between a perturbed design and the reference design. When the number of parameters becomes large, however, it becomes increasingly difficult to isolate a sizable subset without significantly relaxing the similarity requirement across all parameters. While the framework accommodates a large data set in computing model correlation, it is still inappropriate to include correlation information for all available aircraft when the design objective falls within a particular class. For instance, it would be rather misleading to include data from recreational aircraft when computing model correlation if the design objective is a long range military transport aircraft. The models may be correlated differently for the recreational aircraft compared to the military transport aircraft, given that the scale on which certain model physics act is different across the two aircraft types. Thus, it is important to develop a robust and automatic way to identify, out of all available data, a suitable subset for the specific design problem.

There are two different approaches. The first is to use a continuous weighting function instead of the strict cutoff presented in the 2-D problem. Functions such as logistic, piecewise linear, and exponential can all be used to provide continuous weights based on similarity to the reference aircraft. This approach allows all designs to be relevant in computing model correlation, but those more similar to the reference design will be weighted higher than those less similar. The main challenges involved are the potential scale differences among design variables and the amalgamation of weights across all parameters (perhaps not all parameters are equally important).

The second approach makes use of classification algorithms in machine learning. To implement this approach, the engineer must first hand-pick a training set including both similar and dissimilar designs to the reference aircraft. Afterwards, classification tools such as support vector machines, generalized additive models, neural nets, and trees can be trained to classify the rest of the available data[16].

6.2.3 Alternate Correlation Definitions

In the proposed framework, model correlation is given by Pearson's correlation coefficient. Pearson's correlation is a measure of strength and direction of the linear relationship between two variables. While it is mathematically possible to construct a scenario in which model A's error increases nonlinearly relative to model B's, it can only occur when one of the models is extremely inadequate, in which case its surrogate variance should appropriately reflect its fidelity level and render its contribution to the final surrogate insignificant. Furthermore, by stratifying the available data based on its position, we can capture model correlation piecewise locally across the entire design space. Appropriately sized regions should reflect model physics, which should naturally lead to well-defined model correlations.

Nevertheless, there are other types of correlation that have not been explored in the context of the proposed framework. Spearman's correlation, for instance, measures how well the relationship between two variables can be described by a monotonic function. It is quite possible that a different type of correlation will lead to different region widths, which would imply the highlight of physics on a different scale.

It is also worth mentioning that the formulae for information fusion are borrowed from Winkler[30]. In his paper, Winkler defines a correlation measure based on the proportion of points common to each pair of samples. By this definition, a correlation can never be negative, since minimum correlation happens when no points are common to any pair of samples. In Allaire's work[3], however, model correlation is characterized based on model discrepancy trends, which opens up the possibility of negative correlation. This extension of definition prompts further study on this topic, and we have provided a means for continuation through our analysis of results in the 2-D problem.

6.2.4 Interpolation

We have shown that linear interpolation between correlation matrices from adjacent regions produces continuous and positive-definite correlation trends. However, post-smoothing must be conducted on the fused surrogate before gradient-based optimization can be applied. Analysis of results from the 2-D problem reveals that the region containing all positive-definite covariance matrices with respect to pairwise model correlations may be convex. Furthermore, the geometry of the region appears insensitive to the constituent surrogate variances. If the aforementioned speculations are true and can be generalized, it may be possible to devise an interpolation scheme under the positive-definite constraint (specified in analytically tractable form) that preserves smoothness, thereby removing the need for post-smoothing.

6.2.5 Dynamic Surrogate Refinement

The motivation behind surrogate optimization is to reduce cost of computation. Global surrogates, such as ones produced using the proposed framework, may still exhibit inadequacies if the area near the optimum does not carry many training points. Borrowing ideas from recursive Bayesian estimation, it may be possible to devise algorithms for dynamic surrogate refinement, such that new training points can be introduced to some current state of the constituent surrogates without the need to retrain everything from scratch, when the optimizer is nearby.

6.2.6 Quantifying Uncertainty in Model Correlation

An important source of uncertainty that was not accounted for in the fused surrogate comes from model correlation. We have shown in the demonstration problem that model correlation is sensitive to region width. Given that region widths are chosen subjectively, coupled with potential variability from similarity functions discussed in Section 6.2.2, ignoring uncertainty in model correlation could mean large underestimation in fused variances. One place to start might be to track the variability in model correlation at every point (after interpolation) as we vary region widths across a range of values. The extent to which model correlations change may shed light onto its uncertainty.

6.2.7 Multiobjective Optimization

In this thesis, the proposed framework is demonstrated on problems that can be set up for single-objective optimization. Nevertheless, given that models in engineering design often produce multiple outputs, combinations of models may be used to conduct multi-objective optimization. The naive approach to multi-objective optimization in a multifidelity environment assumes independence among model outputs. However, correlation between model outputs may influence fidelity management strategies, provided that improved confidence in one output could lead to improved confidence in a separate, highly correlated output. Li introduces a method for assigning confidence to models with multiple correlated outputs[24].

Bibliography

[1] N. M. Alexandrov, J. E. Dennis Jr., R. M. Lewis, and V. Torczon. A trust-region framework for managing the use of approximation models in optimization. *Structural Optimization*, 15:16–23, 1998.

[2] N. M. Alexandrov, R. M. Lewis, C. R. Gumbert, L. L. Green, and P. A. Newman. Approximation and model management in aerodynamic optimization with variable-fidelity models. *Journal of Aircraft*, 38:1093–1101, 2001.

[3] Douglas Allaire and Karen Willcox. A mathematical and computational framework for multifidelity design and analysis with computer models. *International Journal for Uncertainty Quantification*, 4:1–20, 2014.

[4] Tara Chklovski. Pointed-tip wings at low reynolds numbers. Technical report, University of Southern California.

[5] S. Choi, J. J. Alonso, I. M. Kroo, and M. Wintzer. Multi-fidelity design optimization of low-boom supersonic jets. *Journal of Aircraft*, 45:106–118, 2008.

[6] Daniel Erik Christensen. Multifidelity methods for multidisciplinary design under uncertainty. Master's thesis, MIT, Department of Aeronautics and Astronautics, August 2012.

[7] H. Chung and J. J. Alonso. Design of a low-boom supersonic business jet using cokriging approximation models. *AIAA Paper*, 2002.

[8] H. Chung and J. J. Alonso. Using gradients to construct cokriging approximation models for high-dimensional design optimization problems. *AIAA Paper*, 2002.

[9] Mark Drela. *TASOPT Summary.* MIT, Department of Aeronautics and Astronautics, March 2010.

[10] Mark Drela. *Transport Aircraft System OPTimization.* MIT, Department of Aeronautics and Astronautics, 2.08 edition, February 2011.

[11] M. Ebden. *Gaussian Processes for Regression: A Quick Introduction*, August 2008.

[12] Carl H. FitzGerald, Charles A. Micchelli, and Allan Pinkus. *Functions that preserve families of positive semidefinite matrices.* Elsevier Science Inc., 1995.

[13] Boeing 737-800. `http://www.flugzeuginfo.net/acdata_php/acdata_7378_en.php`, 2001. [Online; last updated 22-Feb-2015; last accessed 17-May-2015].

[14] Alexander I. J. Forrester, András Sóbester, and Andy Keane. Multi-fidelity optimization via surrogate modelling. *Proceedings of the Royal Society A-Mathematical Physical and Engineering Sciences*, 463:3251–3269, 2007.

[15] M. Goldstein and J. Rougier. Reified bayesian modeling and inference for physical systems. *Journal of Statistical Planning and Inference*, 139:1221–1239, 2006.

[16] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning.* Springer Science+Business Media, LLC, 2009.

[17] D. Huang, T. T. Allen, W. I. Notz, and R. A. Miller. Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 32:369–382, 2006.

[18] Paul Jackson, Kenneth Munson, Lindsay Peacock, and Susan Bushell. *Jane's All the World's Aircraft 2009–2010.* IHS (Global) Limited, 2009.

[19] Andy Keane. Wing optimization using design of experiment, response surface, and data fusion methods. *Journal of Aircraft*, 40:741–750, 2003.

[20] M. C. Kennedy and A. O'Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87:1–13, 2000.

[21] M. C. Kennedy and A. O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63:425–464, 2001.

[22] Rémi Lam. Surrogate modeling based on statistical techniques for multi-fidelity optimization. Master's thesis, MIT, Department of Aeronautics and Astronautics, May 2014.

[23] Rémi Lam, Douglas Allaire, and Karen Willcox. Multifidelity optimization using statistical surrogate modeling for non-hierarchical information sources. *AIAA SciTech*, January 2015.

[24] Wei Li, Wei Chen, Zhen Jiang, Zhenzhou Lu, and Yu Liu. New validation metrics for models with multiple correlated responses. *Reliability Engineering and System Safety*, 127:1–11, 2014.

[25] Andrew March and Karen Willcox. Provably convergent multifidelity optimization algorithm not requiring high-fidelity derivatives. *AIAA Journal*, 50:1079–1089, 2012.

[26] P. Piperni, A. DeBlois, and R. Henderson. Development of a multilevel multidisciplinary-optimization capability for an industrial environment. *AIAA Journal*, 51:2335–2352, 2013.

[27] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning.* MIT Press, 2006.

[28] John A. Rice. *Mathematical Statistics and Data Analysis.* Thompson Higher Education, 2007.

[29] E. Torenbeek. Development and application of a comprehensive, design-sensitive weight prediction method for wing structures of transport category aircraft. Technical report, Delft University of Technology, 1992.

[30] Robert L. Winkler. Combining probability distributions from dependent information sources. *Management Science*, 27:479–488, April 1981.

Appendix A

Reference Aircraft Parameters

This appendix details the input parameters to TASOPT associated with the reference aircraft (737-800) used throughout the thesis. Note that for certain parameters, TASOPT includes a conversion multiplier for user convenience.

A.1 Mission Specification

Variable Name	Description	Default Value	Conversion Multiplier
Range	Flight distance	3000.0 [nmi]	1852.0 [m/nmi]
Npax	Number of passengers	180	
Wpax	Passenger weight (including luggage)	215.0 [lb]	4.45 [N/lb]
altTO	Takeoff/landing altitude	0.0 [ft]	0.3048 [m/ft]
T0TO	Ambient takeoff/landing temperature	288.0 [K]	
altCR	Altitude at beginning of cruise	35000.0 [ft]	0.3048 [m/ft]

A.2 Takeoff and Initial Climb Parameters

Variable Name	Description	Default Value	Conversion Multiplier
clpmax	Max perpendicular lift coefficient of wing $C_{Lmax}/cos(sweep)^2$	2.25	
cdefan	Fan drag of dead engine during free climb $C_{DA_{fan}}/A_{fan}$	0.500	
CDgear	Landing gear drag during climb $C_{DA_{gear}}/S$	0.015	
CDspoiler	Spoiler drag during breaking $C_{DA_{spoiler}}/S$	0.10	
muroll	Rolling resistance (friction) coefficient	0.025	
mubrake	Braking resistance (friction) coefficient	0.35	
hobst	Obstacle height	35.0 [ft]	0.3048 [m/ft]
lBFmax	Max takeoff length	8000.0 [ft]	0.3048 [m/ft]
gtocmin	Min top-of-climb gradient	0.015	
thetaCB	Cutback sight angle (figure 1)	40.0 [deg]	
gammaCB	Prescribed cutback climb angle (figure 1)	3.0 [deg]	
gammaDE1	Prescribed descent angle at top of descent	-3.0 [deg]	
gammaDEn	prescribed descent angle at bottom of descent	-3.0 [deg]	

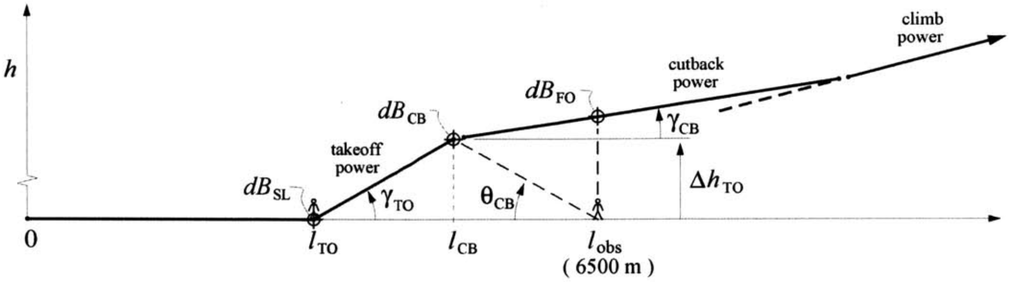

Figure A-1: The visualization serves as a reference for takeoff and climb profile nomenclature. The Figure is taken from [10].

A.3 Sizing-Load Parameters

Variable Name	Description	Default Value	Conversion Multiplier
Nlift	Max vertical load factor for wing bending loads	3.0	
Nland	Max vertical load factor for fuselage bending loads	6.0	
Vne	Never exceed indicated air speed (IAS), for tail loads	280.0 [nm/h]	0.514 [h/nm·m/s]
cabin pressure altitude*	Altitude at which cabin becomes artificially pressurized	7000.0 [ft]	0.3048 [m/ft]

A.4 Cruise-Aero Parameters

Variable Name	Description	Default Value	Conversion Multiplier
CL	Coefficient of lift at cruise	0.58	
Mach	Mach number at cruise	0.78	

A.5 Basic Wing Parameters

Variable Name	Description	Default Value	Conversion Multiplier
sweep	Wing sweep angle	25.33 [deg]	
AR	Overall wing aspect ratio	9.15	
bmax	Max wingspan (for constraint)	117.6 [ft]	0.3048 [m/ft]
lambdas	Inner panel taper ratio c_s/c_o (Figure A-2)	0.70	
lambdat	Outer panel taper ratio c_t/c_o (Figure A-2)	0.25	
iwplan	Wing type: plain cantilever(0), engine mounted(1), or strut braced(2)	1	
ifwcen	Fuel presence (boolean) in center wing box	1	
rWfmax	Usability factor of theoretical max fuel volume	0.90	
fLo	Fuselage lift carryover loss factor (Figure A-3)	-0.3	
fLt	Tip lift rolloff factor (Figure A-3)	-0.05	
yo	Wing centerbox halfspan	71.0 [in]	0.0254 [m/in]
etas	Engine-attach fractional span (Figure A-4)	0.285	

TASOPT models aerodynamic load $\tilde{p}$ in a piecewise-linear fashion. Lift corrections ΔL_o and ΔL_t are applied for fuselage carryover and tip lift rolloff. These values are realized through the adjustment factors fLo (f_{L_o}) and fLt (f_{L_t}). Note that the wing root is located at η_o, and the wing break, where the engine is attached, is located at η_s.

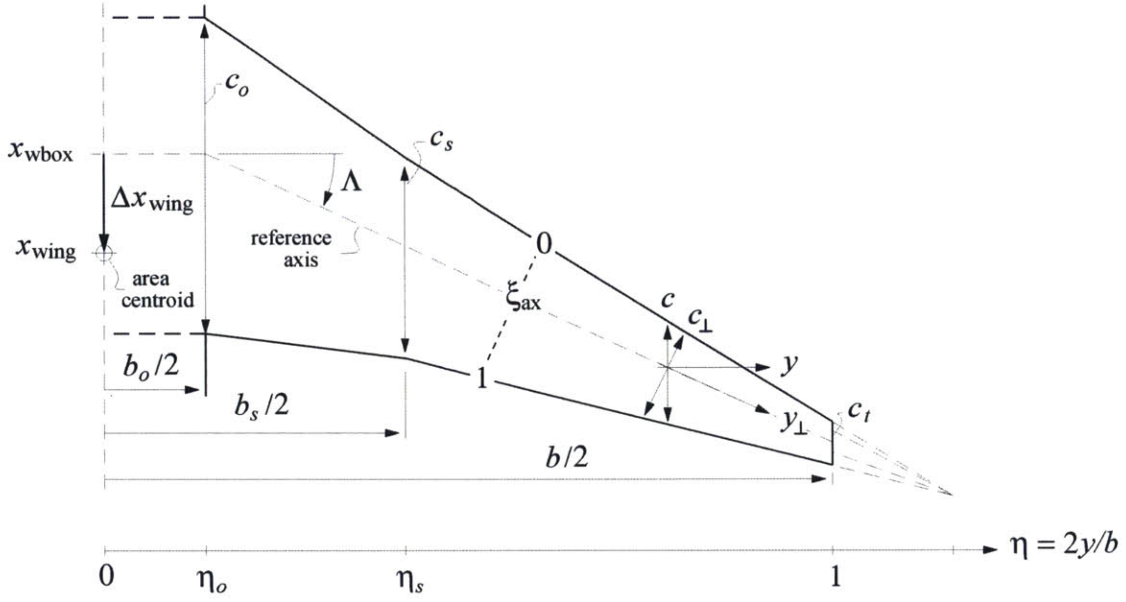

Figure A-2: The diagram serves as a reference for wing geometry variables over a piecewise-linear wing or tail planform. The figure is taken from [10].

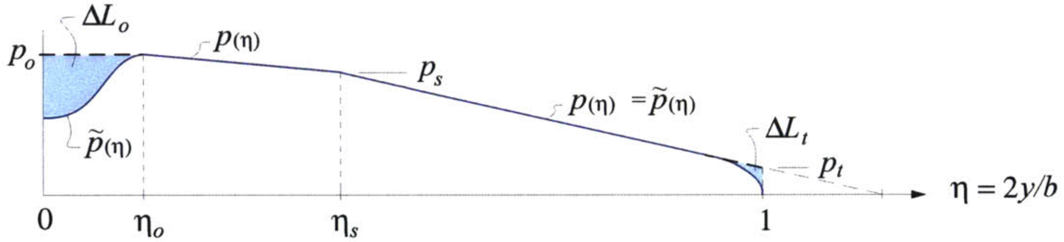

Figure A-3: The solid blue curve shows the piecewise-linear aerodynamic loads across the half-wing. The shaded areas represent the aerodynamic load corrections. The figure is taken from [10].

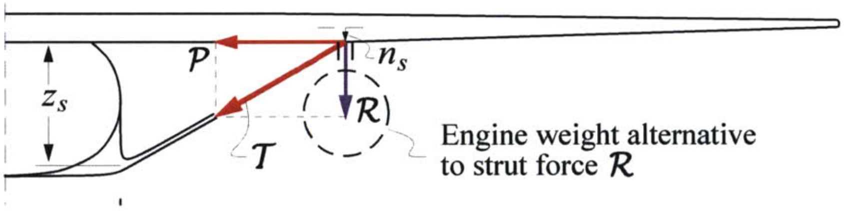

Figure A-4: The diagram shows a cross-section of the wing (cut along the axis of the fuselage), and serves as a reference for engine mounted wing geometry variables. The figure is taken from [10].

A.6 Tail Downward Loading Parameter

Variable Name	Description	Default Value	Conversion Multiplier
CLh/CLmax	Used for tail sizing and wing positioning	-0.5	

A.7 Wing Spanwise Lift and Moment Distributions over Mission

Takeoff, initial climb			
Variable Name	Description	Default Value	Conversion Multiplier
rcls	Wing break/root cl ratio c_{ls}/c_{lo}	1.1	
rclt	Wing tip/root cl ratio c_{lt}/c_{lo}	0.6	
cmpo	Wing profile cm (pitching moment coefficient) at η_o wing root location (Figure A-2)	-0.20	
cmps	Wing profile cm (pitching moment coefficient) at η_s wing break location (Figure A-2)	-0.20	
cmpt	Wing profile cm (pitching moment coefficient) at η_t wing tip location	-0.02	

Clean climb, cruise, descent			
Variable Name	Description	Default Value	Conversion Multiplier
rcls	Wing break/root cl ratio c_{ls}/c_{lo}	1.238	
rclt	Wing tip/root cl ratio c_{lt}/c_{lo}	0.90	
cmpo	Wing profile cm (pitching moment coefficient) at η_o wing root location (Figure A-2)	-0.06	
cmps	Wing profile cm (pitching moment coefficient) at η_s wing break location (Figure A-2)	-0.06	
cmpt	Wing profile cm (pitching moment coefficient) at η_t wing tip location	-0.06	

Landing			
Variable Name	Description	Default Value	Conversion Multiplier
rcls	Wing break/root cl ratio c_{ls}/c_{lo}	1.1	
rclt	Wing tip/root cl ratio c_{lt}/c_{lo}	0.5	
cmpo	Wing profile cm (pitching moment coefficient) at η_o wing root location (Figure A-2)	-0.35	
cmps	Wing profile cm (pitching moment coefficient) at η_s wing break location (Figure A-2)	-0.35	
cmpt	Wing profile cm (pitching moment coefficient) at η_t wing tip location	-0.02	

A.8 Wing and Tail Structural Box Parameters

Variable Name	Description	Default Value	Conversion Multiplier
wbox	Box width/c (Figure A-5)	0.50	
hboxo	Box height/c at wing root (Figure A-5)	0.19	
hboxs	Box height/c at wing break/tip (Figure A-5)	0.109	

Variable Name	Description	Default Value	Conversion Multiplier
rh	Web-height/hbox ratio (Figure A-5)	0.75	
Xaxis	Spar box axis x/c location	0.40	
hstrut	strut t/c (only for strut-braced configurations)	0.15	

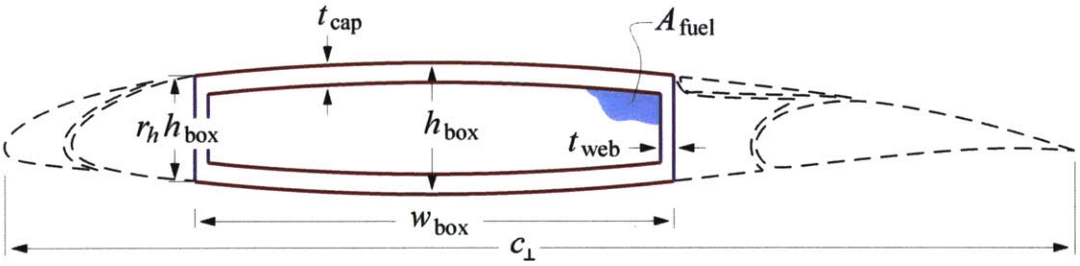

Figure A-5: The diagram shows a spanwise cross-section (perpendicular to the spar axis) of an abstract wing, and serves as a reference for wing/tail box geometry variables. The figure is taken from [10].

A.9 Weight Fractions of Flight Surfaces and Secondary Wing Components (Relative to Wing Box)

Variable Name	Description	Default Value	Conversion Multiplier
fflap	Flaps, flap mounts and actuators weight fraction	0.200	
fslat	Slats, slat mounts and actuators weight fraction	0.100	
faile	Ailerons, aileron mounts and actuators weight fraction	0.040	
flete	Leading and trailing edge weight fraction	0.100	
fribs	Ribs, local stiffeners and reinforcements	0.150	
fspoi	Spoilers, spoiler mounts and actuators	0.020	
fwatt	Wing attachment hardware weight fraction	0.030	

A.10 Horizontal and Vertical Tail (HT, VT) Parameters

Variable Name	Description	Default Value	Conversion Multiplier
iHTsize	Set horizontal tail area (Sh) via Vh(1) or CLhCGfwd at max-forward center-of-gravity (CG) during landing(2)	2	

Variable Name	Description	Default Value	Conversion Multiplier
Vh	HT volume coefficient (used if iHTsize = 1) (Equation 301 in TASOPT documentation[10])	1.45	
CLhCGfwd	HT CL at forward CG trim (used if iHTsize = 2)	-0.70	
iVTsize	Set vertical tail area (Sv) via Vv(1) or CLveout on engine-out(2)	2	
Vv	VT volume coefficient (used if iVTsize = 1) (Equation 302 in TASOPT documentation[10])	0.10	
CLveout	VT CL at engine-out trim (used if iVTsize = 2)	0.5	
iwxmove	Wing position specifier (fixed(0), move to get CLh = CLhspec in cruise(1), or move to get min static margin = SMmin(2))	0	
CLhspec	Specified HT lift coefficient at never-exceed dynamic pressure	-0.02	
SMmin	Minimum static margin (distance between the center of gravity and the neutral point, or aerodynamic center, of the aircraft)	0.05	
dCLh/dCL	HT lift-curve slope factor for neutral point calculation (Equation 292 in TASOPT documentation[10])	0.36	
deps/da	Downwash factor at tail	0.60	

Variable Name	Description	Default Value	Conversion Multiplier
dCLn/dCL	Nacelle lift-slope ratio for neutral point calculation	0.70	
dCLn/da	Nacelle lift-curve slope for neutral point calculation	3.8	
ARh	HT aspect ratio	6.29	
ARv	VT aspect ratio	1.95	
lambdah	HT taper ratio	0.25	
lambdav	VT taper ratio	0.30	
sweeph	HT sweep	36.1 [deg]	
sweepv	VT sweep	21.5 [deg]	
yoh	HT support y location	0.0	
yov	VT support z location	0.0	
fCDhcen	CDhtail (HT drag coefficient) contribution factor for center part $0 < y <$ yoh	0.1	
CLhmax	HT max +/-CL at Vmn (maneuvering velocity), for HT structural sizing	2.0	
CLvmax	VT max +/-CL at Vmn (maneuvering velocity), for VT structural sizing	2.6	
fhadd	HT added-weight fraction (e.g. ribs, LE, elevator)	0.30	
fvadd	VT added-weight fraction (e.g. ribs, LE, rudder)	0.40	
wboxh	HT box width/chord	0.50	
wboxv	VT box width/chord	0.50	
hboxh	HT box height/chord (airfoil t/c)	0.14	
hboxv	VT box height/chord (airfoil t/c)	0.14	

Variable Name	Description	Default Value	Conversion Multiplier
rhh	HT web-height/hbox ratio	0.75	
rhv	VT web-height/hbox ratio	0.75	
nvtail	Number of vertical tails	1	

A.11 Cabin and Fuselage Geometry Parameters

Variable Name	Description	Default Value	Conversion Multiplier
Rfuse	Main bubble radius	77.0 [in]	0.0254 [m/in]
dRfuse	Downward shift of lower bubbles	15.0 [in]	0.0254 [m/in]
wbd	Bubble center y offset	0.0 [in]	0.0254 [m/in]
hfloor	Depth of floor beams	5.0 [in]	0.0254 [m/in]
anose	Nose radius	1.65 [m]	
btail	Tail radius	2.0 [m]	
xnose	Nose tip location	0.0 [ft]	0.3048 [m/ft]
xend	Tail tip location	124.0 [ft]	0.3048 [m/ft]
xblend1	Start of cylindrical section	20.0 [ft]	0.3048 [m/ft]
xblend2	End of cylindrical section	97.0 [ft]	0.3048 [m/ft]
xshell1	Front of pressure shell (center of nose ellipse)	17.0 [ft]	0.3048 [m/ft]
xshell2	End of pressure shell (center of end bulkhead)	102.0 [ft]	0.3048 [m/ft]
xconend	End of tailcone primary structure (distinct from aerodynamic endpoint)	117.0 [ft]	0.3048 [m/ft]
xwbox	Wing box location	57.0 [ft]	0.3048 [m/ft]

Variable Name	Description	Default Value	Conversion Multiplier
xhtail	HT box location	114.5 [ft]	0.3048 [m/ft]
xvtail	VT box location	110.0 [ft]	0.3048 [m/ft]
zwing	Wing box z location	-5.5 [ft]	0.3048 [m/ft]
zhtail	HT box z location	0.0 [ft]	0.3048 [m/ft]
iengloc	Engine location (wing(1) or fuselage(2))	1	
xeng	Engine x location	52.0 [ft]	0.3048 [m/ft]
yeng	Lateral distance of the outermost engine from the centerline	16.0 [ft]	0.3048 [m/ft]
neng	Number of engines	2	
lambdac	Tailcone taper ratio	0.3	
fstring	Stringer weight ratio $W_{stringer}/W_{skin}$	0.35	
fframe	Frame weight ratio W_{frame}/W_{skin}	0.25	
fadd	Added wing weight (flaps+slats+ailerons+ leading/trailing edge+ribs+spoilers+wing attachment hardware) ratio W_{add}/W_{skin}	0.20	
Wfix	Added fixed weight (pilots, cockpit)	3000.0 [lb]	4.45 [N/lb]
xfix	Location of Wfix's center of gravity	7.0 [ft]	0.3048 [m/ft]
W'window	Window weight per unit length $W_{window}/length$	145.0 [m]	3.0 [N/m]
W"insul	Insulation and inside shell weight per unit area $W_{insul}/area$	22.0 $[N/m^2]$	
W"floor	Floor planking (not beams) $W_{floor}/area$	60.0 $[N/m^2]$	

Variable Name	Description	Default Value	Conversion Multiplier
rMh	Inertial-relief factor for HT load on fuselage (for modeling unconstrained structures)	0.4	
rMv	Inertial-relief factor for HT load on fuselage (for modeling unconstrained structures)	0.7	
ifclose	Fuselage tapering option (to a point(0) or to an edge(1))	0	
CMVf1	Fuselage moment volume derivative $d(M_{fuse})/dC_L$	2390.0 [in]	0.0254 [m/in]
CLMf1	CL where $Mfuse = 0$	0.185	
fduo	Fuselage velocity overspeed at wing root	0.018	
fdus	Fuselage velocity overspeed at wing break	0.014	
fdut	Fuselage velocity overspeed at wing tip	0.0045	

A.12 Power Systems and Landing Gear Locations and Weight Fractions

Variable Name	Description	Default Value	Conversion Multiplier
xhpesys	Hyd/pneu/ele system location	62.0 [ft]	0.3048 [m/ft]
xlgnose	Nose landing gear location	14.0 [ft]	0.3048 [m/ft]

Variable Name	Description	Default Value	Conversion Multiplier
dxlgmain	Main landing gear offset behind wing lift centroid	1.0 [ft]	0.3048 [m/ft]
fhpesys	Hyd/pneu/ele system weight ratio (relative to total weight) W_{hpesys}/W_{MTO}	0.010	
flgnose	Nose landing gear weight ratio (relative to total weight) W_{lgnose}/W_{MTO}	0.011	
flgmain	Main landing gear weight ratio (relative to total weight) W_{lgmain}/W_{MTO}	0.044	

A.13 Other Added-Weight Fractions

Variable Name	Description	Default Value	Conversion Multiplier
xapu	Auxiliary power unit (APU) location	120.0 [ft]	0.3048 [m/ft]
fapu	APU weight fraction (relative to payload weight) W_{apu}/W_{pay}	0.035	
fseat	Seat weight fraction (relative to payload weight) W_{seat}/W_{pay}	0.10	
fpadd	Other payload-proportional fraction *e.g. flight attendants, food, galleys, toilets, luggage compartments and furnishings, doors, lighting, air conditioning systems, in-flight entertainment systems, etc.* (relative to payload weight) W_{padd}/W_{pay}	0.35	

Variable Name	Description	Default Value	Conversion Multiplier
feadd	Engine accessories, fuel system fraction (relative to bare engine weight) W_{eadd}/W_{bare}	0.10	
fpylon	Engine pylon weight fraction $W_{pylon}/(W_{ebare} + W_{eadd} + W_{nacelle})$	0.10	
freserve	Reserve fuel fraction (relative to fuel burn) $W_{freserve}/W_{burn}$	0.200	

A.14 Allowable Stresses at Sizing Cases

Variable Name	Description	Default Value	Conversion Multiplier
sigfac	Multiplier on all stress values below	1.0	
sigskin	Allowable fuselage pressurization skin stress	15000.0 [Pa]	0.000145 [psi/Pa]
sigbend	Allowable fuselage bending skin+stringer stress	30000.0 [Pa]	0.000145 [psi/Pa]
sigcap	Allowable wing, tail bending caps stress	30000.0 [Pa]	0.000145 [psi/Pa]
tauweb	Allowable wing, tail shear webs stress	20000.0 [Pa]	0.000145 [psi/Pa]
sigstrut	Allowable strut stress	30000.0 [Pa]	0.000145 [psi/Pa]

A.15 Fuselage Shell Modulus Ratio, for Bending Material Sizing

Variable Name	Description	Default Value	Conversion Multiplier
rEshell	Fuselage shell Young's modulus ratio E_{bend}/E_{skin}	1.0	

A.16 Moduli, for Strut-Induced Buckling Load Estimation

Variable Name	Description	Default Value	Conversion Multiplier
Ecap	Wing sparcap Young's modulus	10.0e6 [Pa]	0.000145 [psi/Pa]
Estrut	Strut Young's modulus	10.0e6 [Pa]	0.000145 [psi/Pa]

A.17 Structural Material Densities

Variable Name	Description	Default Value	Conversion Multiplier
rhoskin	Fuselage skin density	2700.0 $[kg/m^3]$	
rhobend	Fuselage bending stringers density	2700.0 $[kg/m^3]$	
rhocap	Wing, tail bending caps density	2700.0 $[kg/m^3]$	
rhoweb	Wing, tail shear webs density	2700.0 $[kg/m^3]$	
rhostrut	Strut density	2700.0 $[kg/m^3]$	

A.18 Database for Wing Profile Drag in Transonic Cruise, High Climb

Variable Name	Description	Default Value	Conversion Multiplier
cdfw	Wing profile 2D friction drag coefficient for low speed (takeoff, initial climb)	0.0085	
cdpw	Wing profile 2D pressure drag coefficient	0.0035	
Rerefw	Wing profile reference Reynold's number	20.0e6	
cdft	Tail profile 2D friction drag coefficient	0.0060	
cdpt	Tail profile 2D pressure drag coefficient	0.0030	
Rereft	Tail profile reference Reynold's number	10.0e6	
aRexp	Exponent for Reynold's number scaling: $CD = cd * (Re/Re_{ref})^{aRexp}$	-0.15	
fexcdw	Wing excrescence drag factor	1.08	
fexcdt	Tail excrescence drag factor	1.08	
fexcdf	Fuselage excrescence drag factor	1.10	
fBLIw	Fraction of wing boundary layer kinetic energy defect ingested	0.0	
fBLIf	Fraction of fuselage boundary layer kinetic energy defect ingested	0.0	
iBLIc	0 = core in clean flow, 1 = core ingests kinetic energy defect (if any)	0	

A.19 Fuel Parameters

Variable Name	Description	Default Value	Conversion Multiplier
ifuel (methane)	N/A (Not provided in TASOPT documentation[10])	11	
ifuel (ethane)	N/A (Not provided in TASOPT documentation[10])	12	
ifuel (propane)	N/A (Not provided in TASOPT documentation[10])	13	
ifuel (butane)	N/A (Not provided in TASOPT documentation[10])	14	
ifuel (octane)	N/A (Not provided in TASOPT documentation[10])	18	
ifuel (kerosene)	N/A (Not provided in TASOPT documentation[10])	24	
rhofuel (methane)	Fuel density	423.0 $[kg/m^3]$	
rhofuel (ethane)	Fuel density	547.0 $[kg/m^3]$	
rhofuel (propane)	Fuel density	582.0 $[kg/m^3]$	
rhofuel (butane)	Fuel density	600.0 $[kg/m^3]$	
rhofuel (octane)	Fuel density	700.0 $[kg/m^3]$	
rhofuel (kerosene)	Fuel density	817.0 $[kg/m^3]$	

A.20 Engine Temperatures

Variable Name	Description	Default Value	Conversion Multiplier
Tmetal	Turbine metal temperature	1222.0 [K]	
Tt4TO	Turbine inlet total temperature for takeoff	1833.0 [K]	
fTt4CL1	Turbine inlet total temperature fraction for initial climb	0.2	
fTt4CLn	Turbine inlet total temperature fraction for for climb	0.2	
Tt4CR	Turbine inlet total temperature for cruise	1591.5 [K]	

A.21 Turbine Cooling Parameters

Variable Name	Description	Default Value	Conversion Multiplier
dTstrk	Hot-streak temperature allowance	200.0 [K]	
Mtexit	Turbine blade row exit Mach, for temperature drops	1.0	
StA	Area-weighted effective Stanton number	0.09	
efilm	Blade-to-cooling flow heat transfer efficiency	0.72	
tfilm	Cooling-film effectiveness factor	0.32	
M4a	Mach number at start of cooling-air mixing zone	0.9	

Variable Name	Description	Default Value	Conversion Multiplier
ruc	Velocity ratio of exiting cooling air u_{cool}/u_{edge}	0.30	

A.22 Design Pressure Ratios, Efficiencies, etc.

Variable Name	Description	Default Value	Conversion Multiplier
OPR	Overall design pressure ratio	26.2	
pihc	High pressure compressor (HPC) pressure ratio	9.0	
FPR	Fan pressure ratio	1.61	
pid	Diffuser pressure ratio	0.995	
pib	Burner pressure ratio	0.946	
pifn	Fan nozzle pressure ratio	0.9937	
pitn	Core nozzle pressure ratio	0.9919	
epolf	Fan polytropic efficiency at $FPR = FPR_o$	0.923	
epollc	Low pressure compressor (LPC) polytropic efficiency	0.937	
epolhc	High pressure compressor (HPC) polytropic efficiency	0.904	
epolht	High pressure turbine polytropic efficiency	0.871	
epollt	Low pressure turbine polytropic efficiency	0.876	
FPRo	Fan efficiency function constants	1.61	

Variable Name	Description	Default Value	Conversion Multiplier
Kepf	Correction factor for epolf $epolf_{actual} = epolf + K_{epf} * (FPR - FPR_o)$	-0.47	
BPR	Bypass ratio	5.5	
Gearf	LPC/fan speed ratio	1.0	
HTRf	Fan hub/tip ratio	0.30	
HTRlc	LPC hub/tip ratio	0.60	
HTRhc	HPC hub/tip ratio	0.80	
M2	Fan-face Mach number	0.60	
M25	HPC-face Mach number	0.60	

A.23 Fan Nozzle Area Factors Relative to Cruise Design Area

Variable Name	Description	Default Value	Conversion Multiplier
static	Fan nozzle area factor during static state	1.0	
rotation takeoff	Fan nozzle area factor during rotation and takeoff	1.0	
cutback	Fan nozzle area factor during cutback	1.0	
climb1	Fan nozzle area factor during initial climb	1.0	
climbn	Fan nozzle area factor during final climb	1.0	
descent1	Fan nozzle area factor during initial descent	1.0	

Variable Name	Description	Default Value	Conversion Multiplier
descentn	Fan nozzle area factor during final descent	1.0	

A.24 Core Nozzle Area Factors Relative to Cruise Design Area

Variable Name	Description	Default Value	Conversion Multiplier
static	Core nozzle area factor during static state	1.0	
rotation takeoff	Core nozzle area factor during rotation and takeoff	1.0	
cutback	Core nozzle area factor during cutback	1.0	
climb1	Core nozzle area factor during initial climb	1.0	
climbn	Core nozzle area factor during final climb	1.0	
descent1	Core nozzle area factor during initial descent	1.0	
descentn	Core nozzle area factor during final descent	1.0	

A.25 Nacelle Drag Related Variables

Variable Name	Description	Default Value	Conversion Multiplier
rSnace	Nacelle + pylon wetted area/fan area S_{nace}/A_{fan}	16.0	
rVnace	Nacelle local/freestream velocity ratio	1.0	

A.26 Engine Weight Model

Variable Name	Description	Default Value	Conversion Multiplier
iengwgt	Engine weight model selection: McDonnell Douglas original model(0), NF's new model (basic tech)(1), NF's new model (advanced tech)(2); Geared if $Gearf \leq 1$, ungeared of $Gearf = 1$	1	