Predicting Conceptual Aircraft Design Parameters Using Gaussian Process Regressions on Historical Data

Maxfield Arnson [*], Rawan Aljaber [†], and Gokcin Cinar [‡§]

Abstract

This paper presents a novel methodology for predicting key aircraft design parameters using Gaussian Process Regressions (GPR) applied to a comprehensive, open-source database of over 400 aircraft and 200 engines. The database, made freely available through the Future Aircraft Sizing Tool (FAST), enables aircraft designers to apply detailed historical data in early-stage conceptual design and provides full visibility into the data underlying the regressions—unlike traditional regression models, where the training data and model fit are not often disclosed. The non-parametric GPR models developed in this work allow for flexible input configurations, improving the accuracy of predicting critical parameters, such as operating empty weight and uninstalled engine weight, compared to established regressions from the literature. By incorporating a broader range of inputs, these models reduce prediction errors and provide tighter error distributions, leading to more reliable estimates during early design phases. This paper outlines a methodology for adapting conventional aircraft data to explore hybrid-electric and fully electric aircraft designs, ensuring that historical data can be leveraged effectively for novel propulsion systems. Additionally, the flexibility of the GPR framework allows users to create their own regressions and update predictions as new data becomes available, making it a useful tool for researchers working with their own datasets.

Nomenclature

APM	=	Airport Planning Manual
AR	=	Aspect Ratio
ARD	=	Automatic Relevance Determination
BPR	=	Bypass Ratio
CAA	=	Civil Aviation Authority (United Kingdom)
EASA	=	European Union Aviation Safety Agency
EIS	=	Entry-into-Service

[*]Ph.D. Candidate, Department of Aerospace Engineering, University of Michigan, Ann Arbor, Michigan 48109, AIAA Student Member

[†]Graduate Student, Department of Aerospace Engineering, University of Michigan, Ann Arbor, Michigan 48109, AIAA Student Member

[‡]Assistant Professor, Department of Aerospace Engineering, University of Michigan, Ann Arbor, Michigan 48109, AIAA Senior Member

[§]Corresponding author: cinar@umich.edu

EPFD	=	Electrified Powertrain Flight Demonstration
FAA	=	Federal Aviation Administration
FAST	=	Future Aircraft Sizing Tool
GPR/GPM	=	Gaussian Process Regression/Model
M	=	Mach Number
MTOW	=	Maximum Takeoff Weight
OEW	=	Operating Empty Weight
TCDS	=	Type Certificate Data Sheet
T	=	Thrust
$\frac{T}{W}$	=	Thrust-to-Weight ratio
UEW	=	Uninstalled Engine Weight
W	=	Weight
$\frac{W}{S}$	=	Wing Loading

1 Introduction

Conceptual aircraft aimed at sustainably advancing the aviation industry currently lack the streamlined design processes that traditional aircraft benefit from. Typically, modern aircraft design begins with an analytical approach based on historical regressions, drawing from extensive aircraft databases. Major aerospace companies, such as Boeing or Airbus, maintain their own proprietary databases. Popular aircraft design methods, like those by Roskam [1] and Raymer [2], also include regressions intended for conceptual design, though they primarily focus on conventional (i.e. tube-and-wing, Jet-A powered) aircraft. Aircraft design software such as the flight optimization system (FLOPS) [3], or Aviary [4] (which is built upon FLOPS), also utilize historical data for conceptual design. These regressions are used to estimate key performance metrics, such as component weights and geometries, in the early stages of aircraft design. This approach may be refined with later stage regressions, eventually enabling more detailed analyses. However, conventional regression models are limited when it comes to novel aircraft concepts, such as those employing fully or partially electric (i.e., electrified) propulsion systems.

The absence of historical data for such unconventional designs presents an early-stage design challenge; however, recent studies have shown that by revisiting definitions of weight ratios, historical data can provide insights into the feasibility of alternative aircraft as well [5, 6]. Thus, even for novel systems, historical data still holds value in conceptual design. By developing more granular regressions, it is possible to integrate emerging trends, such as new weight and performance metrics, into conventional methods, thereby retaining the decades of experience embedded in historical data. This work provides the foundation for such an approach by leveraging historical regressions and adapting them for novel aircraft concepts, especially those which bridge the gap between contemporary and advanced designs.

This paper aims to provide an up-to-date, open-source database on conventional aircraft that can be easily accessed, manipulated, and analyzed by aircraft designers of all experience levels. The Future Aircraft Sizing Tool (FAST) uses this database for the statistical and regression analysis of aircraft design parameters [7–9]. The methodology for creating and tuning the regressions used in FAST is fully documented, and the models are designed to be adaptable to novel aircraft designs. The results show that the GPRs offer improved accuracy over historical baselines and are well-suited for electrified aircraft design. Section 2 discusses the methodology for this work, including data collection, regression procedures, hyperparameter tuning, as well as the requisite modifications for regression usage with electrified aircraft concepts. Section 3 shows the performance of FAST's OEW and engine weight regressions against historical references, as well as conceptual designs.

2 Methodology

This section describes how data were collected for use in the database, in addition to any post processing that was performed on the data. Additionally, this section describes the methodology used to develop a probabilistic regression to utilize the data in aircraft sizing. It goes on to explain tuning methods which are essential to the functionality of the regressions. Finally, this section explains how the data are used within the FAST tool to assist in quick conceptual aircraft design, with the goal of improving historical regressions on common parameter estimates in the design process, as well as a methodology for adapting conventional data for use in electrified aircraft design.

2.1 Data Collection

To ensure reliable regressions, a substantial database was required. Many existing databases lack proper source citations and seldom offer the level of specificity needed for the intended analysis. For example, while a database might list general technical specifications for the Airbus A380-800, it often overlooks distinctions between variants, such as the -841, -842, and -861, which use different engines.

Published databases vary in their formality and completeness. For instance, the *Jet Engine Specification Database* [10] is a "blog-style" database that provides many aircraft parameters but lacks a large number of unique aircraft. A more formal source is the database compiled by Jenkinson, Simpkin, and Rhodes [11], used in their aircraft design textbook, *Civil Jet Aircraft Design* [12]. This source is thoroughly-documented, offering more categories such as Aircraft, Engines, Airports, Atmosphere and Airspeeds, although includes less detailed information in each category. Janes' *All the World's Aircraft* [13] book is the most widely recognized database, updated annually. Although a website version of Janes' [14] exists, it contains only a portion of what is offered in the book, placing a paywall on accessibility. Similarly, Eurocontrol offers a free, online *Aircraft Performance Database* [15] with limited information and BADA [16] a rich, International Civil Aviation Organization-recognized database that requires payment for any access. Roux's database handbooks on turbofan and turboprop engines [17, 18] provide detailed engine information, but have not been digitized and require manual data entry to use in aircraft design software. The FAA has also published limited datasets online in the form of spreadsheets [19]. Table 1 summarizes existing aircraft and engine databases and their limitations.

Work towards a large, densely populated database for use in FAST [7] is discussed in concurrent work by Arnson et al. [9], which explores relationships between design parameters, key performance parameters, and their physical or regulatory drivers. The study, however, lacks a formal data collection methodology, which is documented here. The so-called FAST Aerobase: Aircraft and Engine Registry Open-Source Database, consists of over 400 aircraft and over 200 turbofan engines, along with a select number of turbojet engines. Each aircraft in the database is represented by 90 parameters, and each engine by 44 parameters[1]. To ensure accuracy, only parameters directly gathered from reliable sources were recorded into the collection. A

[1]Comprehensive list of parameters available on the FAST repository: `https://github.com/ideas-um/FAST` [7]

post-processing script was created and runs in MATLAB to extract any information required from the raw data. Examples of calculated parameters include Operating Empty Weight to Maximum Takeoff Weight ratio (OEW/MTOW), Thrust-to-Weight ratio, Taper Ratio, and Lift-to-Drag ratio, among others.

A key contribution of the FAST Aerobase is the recording of each variant as a unique vehicle, accounting for performance variations resulting from different engine options and configurations. In other words, each combination of an engine and an "airframe" is recorded as its own entry. This level of granularity distinguishes the database from other publicly available resources, making it particularly valuable for detailed performance analysis and conceptual design.

Data were primarily collected from TCDS published online by aviation authorities such as the FAA, EASA, CAA, and others. Additional information was sourced from aircraft manufacturer materials and airport planning manuals. However, fragmented version histories and proprietary information often resulted in an incomplete database. Inconsistencies arose from aircraft variants, certification sheet disparities, and inconsistent conventions. Despite these challenges, a standardized format was adhered to during data collection.

As an example, the data collection process for the Boeing 737-200 began by locating FAA TCDS A16WE on the FAA's website [20]. A portion of this TCDS is seen in Fig. 1. This document lists every certified variant of the 737, including all versions of the aircraft family and certified engine options. The TCDS commonly reports certification year, aircraft weights, maximum operating passenger limits, altitudes, velocities, and thrusts for each engine option. If the data sheet guides a user to notes or flight manuals, those sources are also investigated, such as the payload range diagram for the Boeing 777-200 shown in Fig. 2 [21]. A diagram like this one can be found in a airport planning manual (APM) or "flight manual", which may also contain information about nuanced differences between aircraft. For example, an APM could include the weight differences between two versions of the same aircraft, one with passenger entertainment centers installed and one without. These manuals often show how the payload-range diagrams change as different engines are installed on the same airframe, or have information on the aircraft geometry (e.g. wing area, aircraft bounding box dimensions, etc.) if CAD files are not available on a manufacturer's website.

II - Model 737-200 (Approved December 21, 1967) Transport Aircraft

Engines: 2 Pratt and Whitney Turbofan Engines JT8D-7, JT8D-7A, JT8D-7B, JT8D-9, JT8D-9A, JT8D-15, JT8D-15A, JT8D-17, and JT8D-17A; Refer to the FAA Approved Airplane Flight Manual for aircraft engine and engine intermix eligibility. (Engine Type Certificate No. E2EA)

Fuel: See NOTE 4 for authorized types of fuel.

Engine Ratings:

	Takeoff static thrust, standard day, sea level conditions (5 min) lb.	Maximum continuous static thrust, standard day, sea level conditions lbs.
JT8D-7, -7A, -7B	14,000	12,600
JT8D-9, -9A	14,500	12,600
JT8D-15, -15A	15,500	13,750
JT8D-17, -17A	16,000	15,200

Figure 1: FAA TCDS A16WE's section describing the 737-200.

Once all available aircraft parameters were recorded, the corresponding engines were investigated. For example, the 737 family has used many engines over the years. As shown in Fig. 1, the -200 alone has used nine different engines, and the 737 family consists of 13 commercial variants (-100 through -900 and the MAX -7 through -10). Every engine that has been used by the 737 family was documented in the database.

Table 1: Examples of existing aircraft/engine databases and their limitations.

Author	Year	Database	Type	Limitations
Roux	2007/2011	Turbofan/Turboprop	Book	Format/Date
Janes	2025	Aircraft	Book/Website	Paywall
FAA	2024	Aircraft	Excel	Scale
Eurocontrol	2025	Aircraft	Website	Scale
Eurocontrol (BADA)	2025	Aircraft	Website	Paywall
Jenkinson, Simpkin, Rhodes	2001	Aircraft/Engine	Book/Website	Scale
Meier	2021	Aircraft/Engine	Website	Incomplete

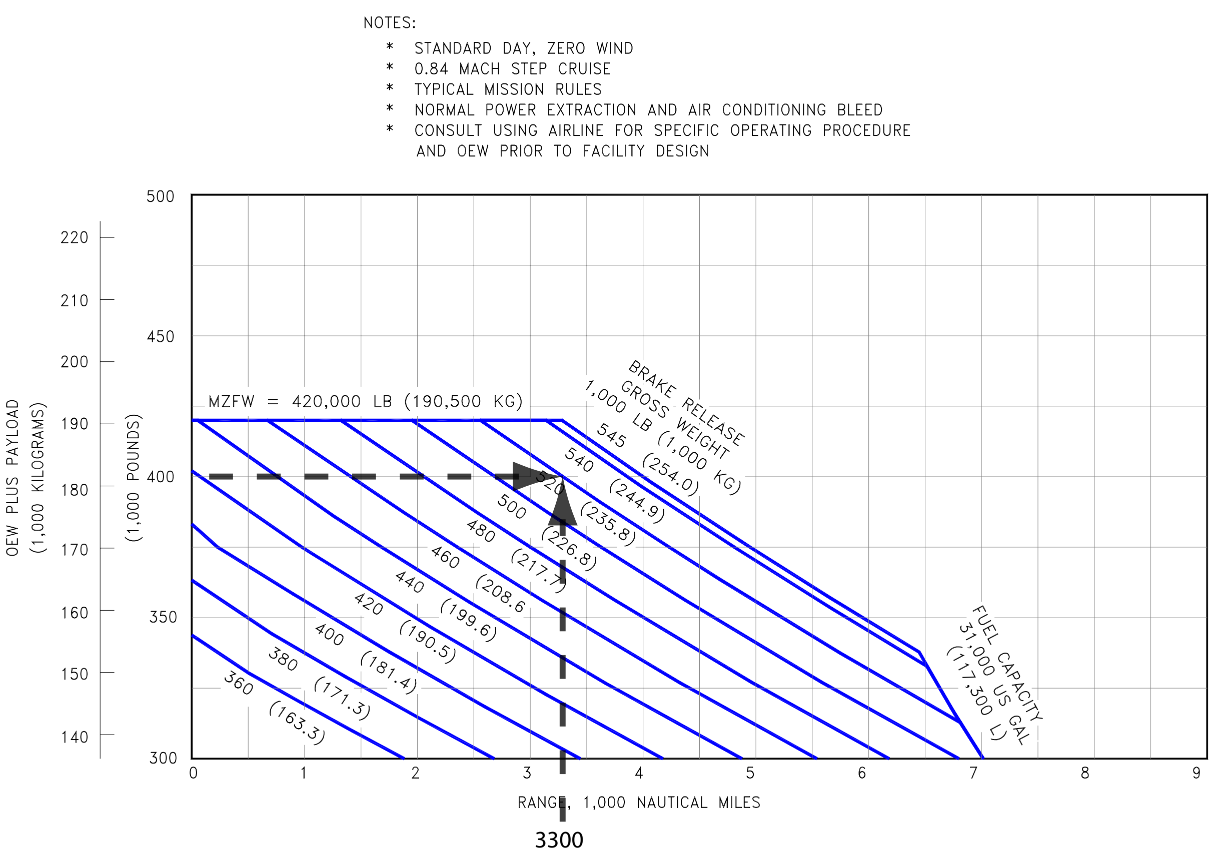

Figure 2: Boeing 777-200 Payload/Range diagram for Mach 0.84 cruise (baseline aircraft).

Engine TCDSs are typically more comprehensive than aircraft TCDSs. Most engine parameters are described in the data sheets. The secondary source of engine information used was Roux's *Turbofan and Turbojet Engines Database Handbook* [17, 18], where information such as fuel consumption, turbine entry temperature, and air mass flow rate can be found. In some cases, additional data were drawn from sources such as a 2008 gas turbine engine data table published in Aviation Weekly [22].

In the event that reliable data could not be found for a specific parameter, no value was recorded. The large collection of aircraft and engine data allowed for robust analysis, even when some entries could not be completed.

2.2 Regression Overview

FAST utilizes regressions which leverage the Aerobase during aircraft sizing. GPRs were chosen for their adaptability to new data and dimensional flexibility, as well as the precedent for their use in aerospace research applications. Chati and Balakrishnan used GPRs to predict MTOW as a function of takeoff roll parameters, trained on data from flight data recorders. The MTOWs were used in a subsequent regression to predict fuel flow in climb out and approach mission segments [23]. Wahler et al. used GPRs as a surrogate model, predicting objective and constraint function values in an aircraft design optimization problem [24]. While Shin et al. also employed machine learning models for aircraft design parameter estimation, their focus was primarily on data imputation for incomplete historical references [25]. This work differentiates itself from previous GPR applications through its novel tuning methods which replace hyperparameter optimization with a generalized data driven method. While perhaps sub-optimal, the autotuning scheme nullifies the need for tuning studies, allowing for the regression framework to perform well unconditionally using any parameters in or derived from the Aerobase.

When using historical data for aircraft design, the number of known point performance metrics may vary between designs. Regressions without predetermined inputs are desirable, because they allow for input flexibility and can make use of all available data. For example, if an aircraft is being designed to carry 150 passengers over 2,000 nautical miles, a good initial guess for the aircraft's maximum takeoff weight (MTOW) will expedite the weight iteration process. If little additional information is available, the regression can use just these two parameters (i.e., payload and range requirements) to predict an initial MTOW. If more information is available, including it as an input will allow the regression to capture more nuanced trends, leading to a more accurate prediction and faster convergence during aircraft sizing.

FAST's regressions are *non-parametric* GPRs, which offer an advantage over a *parametric* GPR by not assuming a predetermined functional output form. This flexibility allows the regression to dynamically learn relationships between aircraft parameters, with the data and the covariance kernel determining the shape of the output space [26]. This capability is especially valuable in aircraft design, where relationships between design parameters are often nonlinear and influenced by multiple factors. Although both parametric and non-parametric GPRs scale computationally with $\mathcal{O}n^3$, the term n represents different quantities in each case. In parametric GPR, n refers to the number of functional coefficients in the output form, while in non-parametric GPR, n refers to the number of data points used for prediction. Given that the database contains hundreds of aircraft, the computational cost of using a non-parametric GPR remains manageable, and can be executed on a laptop computer.

Non-parametric GPR models the relationship between inputs and outputs by assuming that the output data points and any subsets of them are jointly distributed according to a multivariate Gaussian distribution. In FAST, GPR is trained on a set of known input-output pairs from an aircraft database. This process encodes the relationships between the variables, allowing GPR to make probabilistic predictions for unknown outputs. Since the database contains 134 parameters (90 unique to the airframe and 44 unique to the engine) for each entry, FAST first selects the relevant inputs and outputs for each specific regression task. It then creates a targeted subset of variables from the larger set, enabling the regression to focus on the desired relationships.

Non-parametric GPR assumes that the predicted output, given the inputs, follows a Gaussian distribution. As a result, the regression provides a Gaussian distribution for the output, defined by a mean and standard deviation. Equations 1 and 2 present these predicted values. Importantly, GPR does not assume that the input variables themselves follow a normal distribution.

$$\mu = \mu_0 \;\; K\bar{x}, x^* \big(K\bar{x}, \bar{x} \;\; \sigma^2 \times I\big)^{-1} \bar{y} - \mu_0{}^T \tag{1}$$

$$\Sigma = K\bar{x}, \bar{x} - K\bar{x}, x^* \big(K\bar{x}, \bar{x} \;\; \sigma^2 \times I\big)^{-1} K\bar{x}, x^{*T} \tag{2}$$

where:

- μ_0 is the prior mean. It is a default assumption for an output variable if the data proves unhelpful.
- $\bar{x}$ is the set of input variable data for all aircraft used in the regression.
- x^* is the set of input variable data corresponding to the unknown output value. It is the question being asked to the regression.
- σ^2 is the so-called “noise variance." It represents a distrust in the data collected. Analogously, in a physical system this variable captures the effect of noise introduced by a sensor or measuring equipment.
- I is an identity matrix with size $N \times N$, where N is the number of aircraft datasets used in the regression.
- $\bar{y}$ is the set of (known) output variable data for all the aircraft used in the regression.
- Kx_1, x_2 is the covariance kernel. It is a measure of similarity between two sets of data (aircraft in this case), x_1 and x_2. This function and its tuning will be discussed in Section 2.3.

2.3 Covariance Kernel

While non-parametric GPR does not assume a specific functional form for the output, it does require a functional form for the covariance kernel. A covariance kernel is a measure of similarity between two data points, in this case aircraft, with respect to each parameter that defines them. The squared exponential kernel (also called the radial basis function) is widely regarded as a good choice for several reasons. A good kernel should effectively capture appropriate similarity between data points and be computationally efficient [27]. Moreover, the kernel should exhibit decaying similarity as data points become sufficiently different [26, 28].

The squared exponential kernel satisfies these requirements and is adaptable to higher-dimensional spaces while maintaining a relatively low level of computational complexity. The $D-$dimensional form is presented in Equation 3.

$$Kx, x^* = \tau \exp\left(-\gamma \, {}_{i=1}^{D} \frac{x_i - x_i^{*2}}{\ell_i^2}\right) \tag{3}$$

where:

- x and x^* are example data points. Each is an $D-$dimensional vector.
- τ is the output scale factor. It determines the absolute value of the impact that two “similar" data points will have on the posterior distribution.
- γ is the input scale factor. It governs the sensitivity of the kernel to input similarities.
- ℓ_i is the parameter length scale. This takes unique values depending on the parameter being compared. It serves to normalize between parameters so they can be directly summed.

The squared exponential kernel fulfills all requirements for a good choice of covariance kernel. It compares all dimensions of two data points, normalizing each dimension such that all contribute equally. It decays to little contribution as two data points become dissimilar. In effect, this kernel will not modify a prior belief about the regression output if there does not exist data to claim that the prior belief is flawed. It scales computationally with $\mathcal{O}D$, and in practice the dimensionality of a regression is limited to $D \sim 10$ as there are rarely aircraft design parameters with strong relationships to more than 10 other parameters without redundancy.

2.4 Parameter and Hyperparameter Tuning

Thus far, many regression parameters and kernel hyperparameters have been introduced and defined. This section presents studies and discussions on tuning these parameters within the context of regression models. In FAST, the tuning methods are often parameterized based on the available data from the aircraft database, as the regressions are dynamic. Rather than optimizing parameters for a specific regression, tuning studies aim to minimize average errors across multiple regressions.

2.4.1 Prior Mean

The prior mean, shown by μ_0 in Equation 1, represents a prior belief about the output of the regression, which may be adjusted depending on the data’s relevance to the desired output parameter. This prior value can be based on an educated guess, a lower-fidelity numerical model, or a physics-based model. In FAST, since the regression does not have prior knowledge of which dataset corresponds to the output parameter, the prior mean is tuned to the mean of the desired output parameter (as shown in Equation 4). This approach provides the regression with a sense of scale, adjusting it according to the available data without requiring additional models. FAST does allow for an alternative prior to be input into a regression should a user decide they would like to overwrite the average value provided by the database.

$$\mu_0 = \frac{1}{N} \, {}_{i=1}^{N} \bar{y}_i \tag{4}$$

where $\bar{y}$ is defined as it was in Equation 1, the set of output parameter values from all aircraft used in the regression.

2.4.2 Length Scales and Output Scale Factor

The length scale hyperparameter, denoted as ℓ in Equation 3, determines how similar two aircraft parameter values must be for the regression to treat them as “alike." This hyperparameter is highly sensitive to the specific aircraft parameter it applies to, and a distinct length scale is required for each dimension being compared. For example, when predicting operating empty weight (OEW) as a function of maximum takeoff weight (MTOW) and entry-into-service (EIS) year, the regression evaluates the MTOW, EIS, and OEW of aircraft in the database to estimate an appropriate OEW for the given MTOW and EIS. While comparisons between aircraft within 10,000 pounds of the target MTOW may be reasonable, comparisons to aircraft that entered service within 10,000 years of the target EIS would clearly be inappropriate.

There are two primary approaches to tuning the length scale parameter(s). One option is to perform empirical testing or optimization to determine the values that minimize error using a set of validation data. Alternatively, the parameters can be learned using automatic relevance determination (ARD). ARD requires that all available parameters in the database be used for every output, and the length scales are learned by tuning the input variables. Inputs with little relevance to the output (i.e., those showing weak correlation with the output) are naturally assigned near-zero contributions [29].

Since length scales are not universal across different regressions, optimization would require fixing the inputs and outputs of specific regressions, which undermines the flexibility of a dynamic regression model. While ARD could dynamically learn the length scales, this approach would be computationally expensive, negating the benefit of a quick regression call within the aircraft sizing code. However, the two desirable functions that ARD performs, parameter relevance and tuning, can be approximated as follows:

1. When performing a regression, it is assumed that not all input parameters are relevant. Either a select few are chosen based on known physical relationships between aircraft parameters, or a small subset of parameters is used and assumed to be relevant. The former method is employed for regressions with predefined outputs, such as OEW or propulsion system weight. The latter approach is used when the output is set dynamically based on user interaction. When a user inputs information describing an aircraft, these values are used to create a regression that predicts any missing data the user has not provided. In this case, required information such as payload and range is supplemented with optional data to predict the remaining inputs necessary for the sizing iteration.

2. Regarding the tuning of parameters deemed relevant, the goal is not to find the perfect value that minimizes error for a specific regression. Instead, the tuning value should be parameterized based on the data and tailored to the specific input variable. Using the standard deviation of the input variable provides a length scale for the kernel that is both appropriate in scale and intuitive. If two aircraft have input parameter values within one standard deviation of each other, it is reasonable to consider them “similar," as their difference is essentially negligible during the large data processing that the regressions employ.

Equation 5 illustrates how the length scale hyperparameters are tuned.

$$\ell_i^2 = \mathrm{var}\bar{x}_i \tag{5}$$

Since all inputs are normalized by their respective length scales, the kernel output will reflect changes on a consistent order of magnitude. To scale this to the appropriate output variable, an additional tuning parameter, denoted by τ in Equation 3, is required. The exponential term determines the similarity between two data points, while the output scale factor τ dictates the extent to which this similarity influences the adjustments made to the prior belief. Similar to the length scales, τ is tuned to one standard deviation of the output variable, which provides an appropriate scale for the kernel to modify the prior with. This is shown in Equation 6.

$$\tau = \sqrt{\mathrm{var}\bar{y}} \tag{6}$$

2.4.3 Noise Variance

In a physical system, the noise in recorded data may be defined through testing sensors. However, in a simulated system, noise variance is treated as an additional hyperparameter and thus requires tuning as well. Despite being collected from reputable sources, it is ill-advised to assign a small universal value to the noise variance. In addition, the noise should be scaled depending on the regression's output parameter. During data collection, disagreements between sources for parameter values were not uncommon. For example, the range reported in a TCDS may differ from that in an Aircraft Planning Manual. TCDS figures typically reflect values at their extremes (i.e. do-not-exceed speeds or weights), whereas Aircraft Planning Manuals offer detailed diagrams that present multiple specific ranges under various conditions. Typically, these differences were on the order of about 5 to 10 percent. Therefore, it is assumed that the standard deviation of the data noise is 7.5% of the mean of the data for the desired output parameter. Introducing noise variance also prevents against overfitting, as data are not fully trusted to represent the truth.

2.4.4 Input Scale Factor

The input scale factor γ is commonly accepted to be on the order of $\gamma = 0.5$. This parameter is meant to be constant for any regression regardless of the input and output spaces. This was tested empirically by considering multiple regressions which are commonly used in aircraft design. The aircraft database was divided into two sets of data, 90% being used to train the model while 10% was reserved to test the regression output. The data were assigned to validation and training sets randomly. Six common regressions, summarized in Table 2, were then performed using the test data's inputs. This was repeated as the input scale factor γ was varied. Then, a new 90-10% split was taken and the regressions were repeated (100 times) for each γ. All of the errors for each value of γ were averaged for each regression. Additionally, the average of the six regression's error was taken at each gamma to confirm if there exists an optimum value of γ such that the average error across all regressions is minimized. The results from this study are shown in Fig. 3, where it can be shown that for values of $0.2 < \gamma < 5$ the resultant average error across all regressions appears to be somewhat constant. A slight optimum was found at $\gamma = 2.27$, which is the value assumed to universally minimize the error output for any regression in FAST.

Table 2: Regression inputs and outputs for a γ tuning parameter study, shown in Fig. 3

Inputs	Output
MTOW	Aircraft Length
MTOW, Range	OEW
MTOW, Thrust	Engine Weight Fraction
Payload, Range, EIS	MTOW
Payload, Range, MTOW	Lift-to-Drag Ratio
MTOW, EIS	Wing Loading

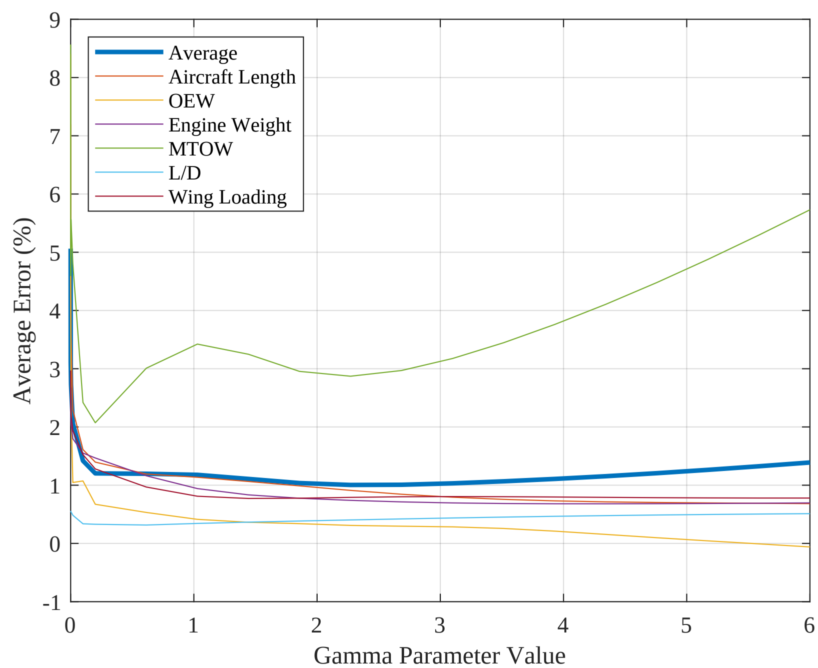

Figure 3: Example regression errors as function of the hyperparameter γ

2.5 Regressions Used in FAST

FAST utilizes several regressions during the aircraft sizing process. Figure 4 outlines FAST's workflow, and highlights the code blocks which call regressions in green, and can be further distinguished with an "R" character. Figure 4 was generated using OpenMDAO [30] as a visual aid, however, FAST is not compatible with the Python-based framework. The initialization block calls regressions to set values for unspecified variables which FAST requires to run. The weight build-up and engine sizing blocks predict OEW and engine weights respectively. These regressions are run inside the sizing iteration, updating their predictions depending on the current iteration's MTOW, thrust, etc. The regression inputs and outputs are as follows:

Initialization: FAST parses user specified values from the following list of parameters:

- Cruise Speed (if unspecified in mission profile)
- Cruise Altitude (if unspecified in mission profile)
- MTOW (initial guess)
- Wing loading
- Lift-to-drag ratio (turbofans only; insufficient turboprop data[2])
- Thrust-to-weight ratio (turbofans) / Power-to-weight (turboprops)
- SLS thrust (turbofans) / SLS power (turboprops)

Then the known values, in combination with required inputs, payload weight and range, are used to predict the unspecified values. The initialization regressions are not updated as the aircraft sizes.

Operating Empty Weight: Described in Sec. 2.5.1 below.

Engine Weight: Predicted using required thrust or power at takeoff, depending on whether the aircraft uses turbofan or turboprop engines.

FAST is an actively developed tool, and is constantly being improved. As more data are collected, the regressions FAST uses are updated. At the time of this work's publication, FAST is under review, and the newest and most accurate regressions are being integrated into the aircraft sizing procedure.

2.5.1 Operating Empty Weight: A New Approach

Although FAST's database consists solely of conventionally powered aircraft, the tool is built to design, analyze, and evaluate advanced aircraft with novel propulsion systems, including electrified aircraft propulsion [8]. Accordingly, when using conventional aircraft data to draw conclusions about hybrid-electric or fully electric concepts, modifications must be made to account for biases inherent in conventional regressions. The following methodology, adapted from Arnson et al. [6], outlines an approach for modifying the inputs and outputs of a conventional OEW regression for use in conceptual design of electrified architectures.

This work follows precedent set by FLOPS [3], whereby engine weight refers to the uninstalled (sometimes called dry) engine weight (UEW), and additional propulsive components like thrust reversers, engine

[2] L/D is predicted using aircraft geometry data, which is sparse for turboprops. In contrast, P/W ratio or SLS power for turboprops are commonly reported in TCDSs or APMs.

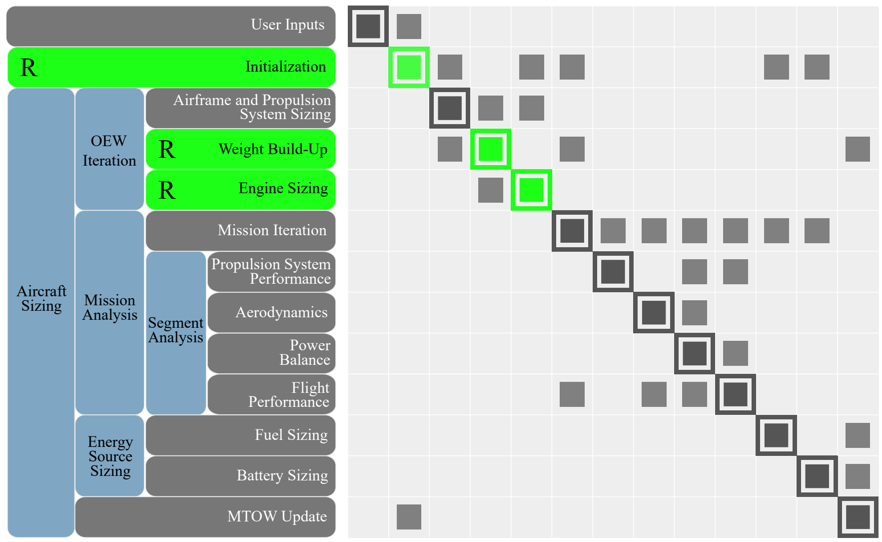

Figure 4: FAST's overarching workflow with regression-using blocks highlighted in green and marked with "R."

controllers, etc. are book-kept to a general "propulsive" group. Nacelles and pylons are book-kept to the airframe (sometimes called structural) group. This scheme is helpful when using historical data because engine manufacturers often report uninstalled weights for their engines, making separate numerical analysis straightforward. Precedent for the electrified aircraft propulsion (EAP) weight group is set in Chau et al. [31], where EAP weights include electric propulsors, power management, thermal, and battery systems.

While OEW typically includes both engine and EAP weights, these are distinguished from other component weights as their sizing is generally more granular due to higher sensitivity to aircraft performance. As an example, landing gear sized to carry a 100 ton aircraft is agnostic of the powerplant aboard the vehicle; relationships driven by conventional aircraft data can be used regardless of the propulsion architecture. Defining a new parameter, which isolates engine and EAP weights, allows the grouping of weights which are much less (if at all) sensitive to propulsion system architecture. The airframe (structural), systems, and operational (ASO) weight is defined in Eq. 7 as

$$\begin{aligned} W_{\text{ASO}} &= W_{\text{MTO}} - W_{\text{Fuel}} - W_{\text{Payload}} - W_{\text{Engines}} - W_{\text{EAP}} \\ &= W_{\text{OE}} - W_{\text{Engines}} - W_{\text{EAP}} \end{aligned} \tag{7}$$

Equation 7 assumes that OEW is defined as $W_{\text{OE}} = W_{\text{MTO}} - W_{\text{Fuel}} - W_{\text{Payload}}$. Because OEW and engine dry weight (EAP weight assumed zero for conventional aircraft) are readily available in the database, the ASO weight is easily added as a database parameter. If the ASO weight is thought of as everything needed to carry the propulsion systems and payload, these terms can be thought of as another weight group, defined in Eq. 8 below as the "burden" weight:

$$\begin{aligned} W_{\text{Burden}} &= W_{\text{Fuel}} \;\; W_{\text{Payload}} \;\; W_{\text{Engines}} \;\; W_{\text{EAP}} \\ &= W_{\text{MTO}} - W_{\text{ASO}} \end{aligned} \tag{8}$$

As maximum takeoff weight is also readily available in the database, the burden weight is easily calculated as well. A regression is then performed predicting ASO using burden weight, maximum takeoff weight, design range, and takeoff performance (thrust for turbofan aircraft or power for turboprops[3]).

FAST uses this procedure for all aircraft sizing, regardless of architecture, including conventional aircraft. This allows for easy handling in the OEW prediction in FAST. The user-input architecture matrices, described in detail in Mokotoff et al. [8], tell FAST which propulsion models to run. These may then output weights for gas turbines (called via the aforementioned engine weight regression), electric motors, fuel, and batteries. In the case that an architecture does not use one of these components, its weight is kept as its initial value (zero) and it does not impact the sizing process. All sized components are added back to ASO to report a final OEW.

This methodology is limited to application for tube and wing aircraft with electrified powertrains, as the component assumptions implicit in the database's ASO weights may not be valid for aircraft utilizing hybrid wing body architectures or hydrogen as a fuel.

3 Results

3.1 Historical Database

The historical database is available online, free of charge, in several formats. Large Excel spreadsheets are available for use with statistical software, in addition to MATLAB data structures which contain the same information. These data structures, which include both engine and aircraft data, are formatted to integrate seamlessly with FAST [7], allowing designers to create their own regressions.

3.2 Regressions

As previously mentioned, the FAST tool is capable of creating a large number of unique regressions using the information in the database. Users are encouraged to design their own input/output spaces, and explore which regressions may be useful in conceptual aircraft design. It is recommended to use parameters which have physical relationships or are highly correlated. Arnson et al. [9] discuss aircraft performance parameters and their key physical drivers in great detail, and the input/output spaces of the FAST regressions were influenced by the study. Of the possible regressions, OEW and engine weight are of the most important, since they are used inside of an aircraft sizing loop. This section compares the FAST regressions to established sources, using the moments of each method's error distribution. The data used for the studies presented in this section come from commercial aircraft utilizing turbofan engines only, and is sourced directly from FAST's database. Appendix 4.1 provides a comprehensive tutorial for recreating the regressions presented in this section. In short, a bootstrapping method is applied to the FAST data. Each aircraft or engine (represented by a set of design variables and performance parameters) is temporarily removed from the database. This modified database is used to train a regression, which then predicts the output (OEW or engine weight) given the requisite inputs. Then the true value of the output is compared to the predicted value, the error recorded, and the aircraft or engine returned to the database. The process is then repeated for all aircraft or engines in the database, ensuring that input/output pairs tested by the regression were not used when training it. This method is also known as "leave one out cross validation," as it ensures that the model is not trained with information it will later be tested on.

[3] OEW data, and therefore ASO weights, are sparse for turboprops. FAST uses a simple linear regression to predict ASO weight for turboprops using a limited dataset.

3.2.1 Operating Empty Weight Regression

As this work discusses, OEW is an important parameter in conceptual aircraft design. Its estimation as a function of other aircraft parameters is essential in FAST during the sizing process. Canonical aircraft design literature, such as Raymer [2], Jenkinson [32], and Roskam [1], offers data-driven regressions to predict OEW. FLOPS [3] includes a data driven OEW buildup as well. Naturally, these established sources are good metrics for which to compare the methodology presented in this work. A multivariate linear regression [33] is also performed as a "control" or benchmark for the FAST regression. This section will compare methodologies, discussing their intended use-case in the aircraft design process, their required inputs, and the moments of their error distributions when used on the FAST Aerobase.

FAST: While canonical sources explicitly predict OEW, FAST predicts ASO weight, consistent with the methodology outlined in this work. It uses burden weight, MTOW, range, and SLS thrust to do so. To compare OEW directly with the other methodologies, FAST adds propulsion weights from the Aerobase back to ASO, resulting in an OEW (only true for conventional configurations as EAP weights are zero).

Multivariate Linear Fit [33]: The multivariate linear regression is intended to define the performance difference between the Aerobase and the GPR methodology described in this work. It is possible that the Aerobase is so expansive that any reasonably good regression method would be able to predict OEW as well as the GPRs. This multivariate linear regression was chosen due to its ease of use in MATLAB, and in this study the input space is identical to the FAST regression: burden weight, MTOW, range, and thrust. This method also predicts ASO weight, adding a known propulsion system weight to produce OEW.

FLOPS [3]: FLOPS was developed at NASA Langley Research Center and has been used for decades for conceptual and preliminary aircraft design. Originally written in Fortran, the code has been adapted into Python and lives on in NASA's Aviary codebase [4], where its aerodynamic drag and structural weight estimation methods are still used today. The weight estimation method comprises dozens of regressions based on historical data. For this study, the method was adapted into MATLAB for use with the FAST Aerobase. The adapted version's assumptions and verification study are included in Appendix 4.2. A clear disadvantage of this methodology is the required assumptions and relatively large input space, however, this work assumes that, for the purposes of comparison with FAST, these parameters are known in the conceptual design stage.

Raymer [2]: Dr. Daniel Raymer is the author of *Aircraft Design, A Conceptual Approach*, and is widely recognized as a leading authority in the field of aircraft design. His book, now in its seventh edition, has been an invaluable resource for both students and practitioners since its first publication in 1989. The book includes three major OEW regressions, each corresponding to a different level of fidelity:

1. A simple relationship between MTOW and OEW, intended for very early conceptual design (source Table 3.1).
2. A mid-fidelity regression, that uses MTOW, aspect ratio, thrust-to-weight ratio (takeoff), wing loading (takeoff), and maximum operating Mach number (source Table 6.1).

3. A high-fidelity component buildup, similar to FLOPS but requiring more information. Using assumed values for this methodology would undermine the benefits of its higher fidelity. While powerful, this level of detail is perhaps better suited for a later stage in the aircraft design process.

The mid-fidelity regression is selected as the most appropriate of Raymer's regressions for comparison with FAST in this work, and it is presented in Eq. 9.

$$W_{\text{OE}}\text{lbm} = 0.869 \times W_{\text{MTO}}^{1\text{ - }0.037}\text{lbm} \times AR^{0.0398} \times (\frac{T}{W})_{\text{MTO}}^{0.1} \times (\frac{W}{S})_{\text{MTO}}^{0.1}(\frac{\text{lbm}}{\text{ft}^2}) \times M_{\text{max}}^{0.05} \tag{9}$$

Jenkinson [12][32]: Dr. Lloyd Jenkinson has authored two books on aircraft design: *Civil Jet Aircraft Design* and *Aircraft Design Projects for Engineering Students*. The former of which offers two methods for predicting OEW: 1) an OEW/MTOW vs MTOW relation, unaccompanied by an equation, and 2) a detailed weight buildup similar to FLOPS and Raymer's high fidelity estimation. *Aircraft Design Projects for Engineering Students* includes an updated version of the first method from *Civil Jet Aircraft Design*, which gives linear relations for OEW as a function of MTOW and engine count. This method is preferred over the detailed weight buildup for the same reasons discussed for Raymer's approach; parameters which are unknown in the early conceptual design stage, such as distances between engines and empennage areas, are required for its use. The OEW versus MTOW and engine count relation used is shown in Eq. 10 below:

$$W_{\text{OE}} = \begin{cases} 0.55 \times W_{\text{MTO}} & N_{\text{Engines}} = 2 \\ 0.47 \times W_{\text{MTO}} & N_{\text{Engines}} > 2 \end{cases} \tag{10}$$

Roskam [1]: Dr. Jan Roskam, founder of Design Analysis and Research Corporation (DARC) and author of several books and over 150 articles on topics such as aircraft design, performance, controls, and dynamics, was globally known for his expertise in aeronautics. His book, *Airplane Design*, first published in 1985 and comprised of eight parts, outlines the aircraft design process. *Part I: Preliminary Sizing of Airplanes*, provides a regression for estimating OEW using MTOW, shown in Eq. 11 below (source Eq. 2.16 and Table 2.15):

$$W_{\text{OE}}\text{lbm} = 10^{-0.0802} \times W_{\text{MTO}}^{-0.0369}\text{lbm} \tag{11}$$

Predicted vs actual plots comparing these methods are illustrated in Fig. 5; these are useful metrics to visually interpret the OEW trends produced by each method. Each regression requires a unique set of input variables to run. For example, FLOPS uses a set of 14 inputs, which is only complete for a subset of the database. As mentioned in Sec. 2.1, not all aircraft have complete entries due to lack of source data, which explains why the methods requiring a smaller input space appear more dense in Fig. 5.

Since the true value of OEW is known for database aircraft, each method's predictions are used to calculate an error according to Eq. 12.

$$\text{Error\%} = \frac{\text{Predicted} - \text{Actual}}{\text{Actual}} \times 100 \tag{12}$$

The moments of the resultant error distributions for each methodology, in addition to the median, are calculated. Table 3 gives a detailed comparison of these error distributions. All methods give reasonable mean errors, under 4%, which suggests that on average all methods are predicting OEW reliably. Additionally, all methods have medians within 1% of their means, indicating low bias. However, FAST has the lowest standard deviation of all the methods, meaning that approximately 68% of the OEW predictions from FAST fall within about 5.05% of the mean (assuming near-normal distribution, discussed below). FAST's mean error is nonzero, but is not large enough to offset the benefit of this low standard deviation. Most error distributions have skewness values between -2 and +2, and proper kurtosis[4] below 7, with the exception of the multivariate linear fit. Following West et al. [34], distributions within these limits are considered not substantially non-normal. FAST's error is more skewed than Raymer and Jenkinson distributions, however shows smaller standard deviations and kurtosis values. This suggests that while FAST may be more likely to overpredict than underpredict, to a higher degree than the Raymer or Jenkinson methods, it is more consistently predicting near the mean, and when it predicts incorrectly, the error is more likely to be smaller in magnitude than that of the Raymer or Jenkinson methodologies. Despite FLOPS outperforming FAST with respect to skewness and kurtosis, FLOPS underperforms in the other error metrics. The multivariate linear fit outperforms FAST only with respect to median, suggesting that a simple linear model, while perhaps trivial to implement, can be significantly improved upon. This work argues that a holistic approach should be taken to determine which methodology is most consistent, and given that FAST has the lowest standard deviation with middle-of-the-pack performance in the other metrics, it is the most consistent and reliable method for predicting OEW.

Table 3: Error metric summary for FAST and reference regressions on OEW.

Method	Sample Size	Mean [%]	Median [%]	Std. Dev. [%]	Skewness	Kurtosis
FAST	183	+0.2941	-0.3017	5.048	+0.2779	3.660
Multivariate Linear Fit	183	+1.761	-0.2494	9.960	+1.767	8.330
FLOPS	92	-3.989	-4.135	8.346	-0.2009	3.037
Raymer	178	-1.074	-1.795	7.426	-0.2223	4.260
Jenkinson	305	-3.887×10^{-4}	+0.9278	8.798	+0.1359	4.572
Roskam	306	+0.2523	-1.176	9.551	+0.7269	3.964

3.2.2 Engine Weight Regression

Uninstalled engine weight (UEW) is a component of the burden weight, which has been shown to be a good predictor of ASO weights. UEW, following FLOPS convention [3], can be thought of as the "engine core," where other propulsive elements, pylons, and nacelles are book-kept to the ASO weight. If a configuration is known to omit some or all additional propulsive components (e.g., fuel systems, thrust reversers), users may provide supplemental weight models to FAST to refine the ASO prediction accordingly. FAST does not currently include any propulsive component models, intending to avoid undermining the purpose of using as few inputs as possible; limited publicly available data would also hinder efforts to create these models. In the early conceptual design stage, UEW may not be known, and can be estimated using a data-driven

[4]Some sources may subtract 3 from kurtosis, such that a normal distribution has kurtosis of 0. "Proper" or "absolute" kurtosis, reported in this work, does not subtract 3 [34].

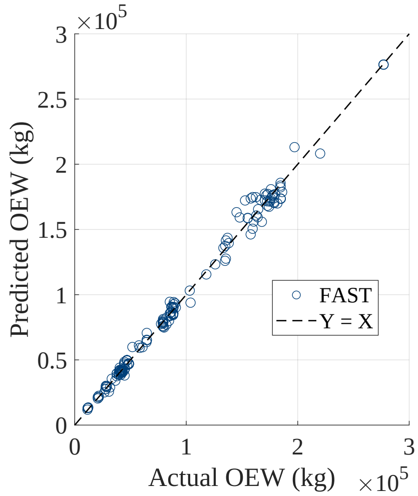

(a) FAST's GPR.

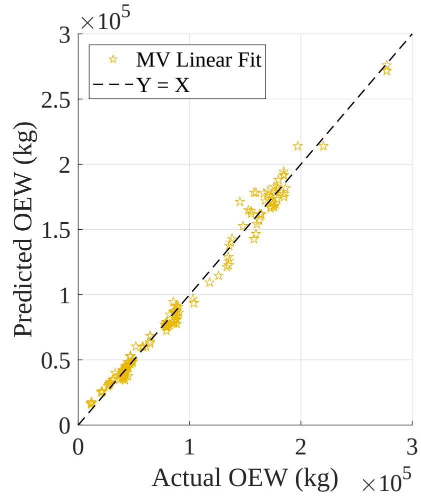

(b) Multivariate linear fit [33].

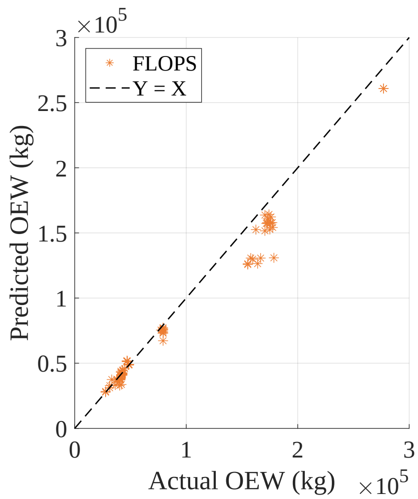

(c) FLOPS buildup [3].

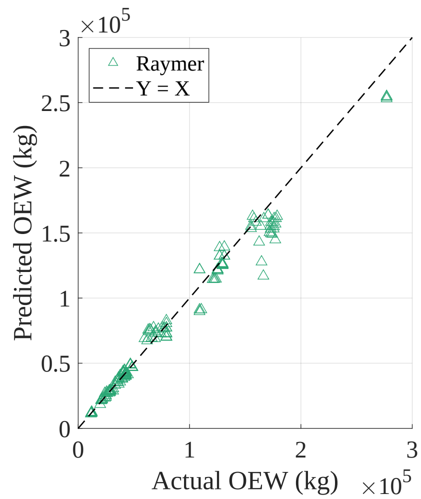

(d) Raymer regression [2].

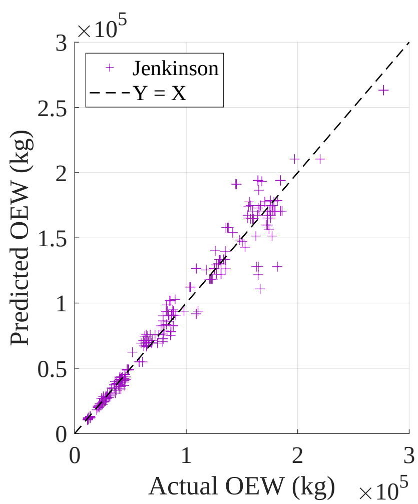

(e) Jenkinson regression [12].

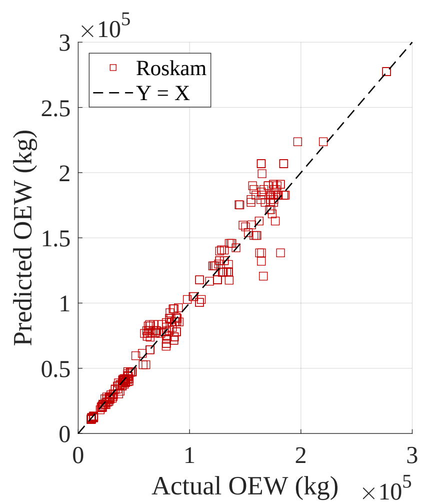

(f) Roskam regression [1].

Figure 5: OEW predictions vs actual for several estimation methods.

model. Previously discussed canonical aircraft design literature typically offers UEW regressions, such as Raymer, Jenkinson, FLOPS, and Svoboda [35] (not previously introduced). As with OEW, a multivariate linear regression is performed as a control. This section describes the UEW prediction methodologies, their required inputs, and the resulting error distributions when applied to the FAST Aerobase.

FAST: FAST uses a GPR that predicts UEW as a function of MTOW and maximum continuous thrust, per engine, at SLS conditions. This value is often reported in engine TCDSs, making it readily available in the Aerobase.

Multivariate Linear Fit [33]: The reasoning for the multivariate linear fit is identical to that of OEW. It is intended to show that good choice of input space and large database is only part of the cause of FAST's improvement over other methods. The input space is identical to the FAST regression: MTOW and thrust (per engine) at SLS conditions.

FLOPS [3]: FLOPS includes a regression for UEW (source Eq. 76). Assuming inlet and nozzle weights have not been independently specified, this regression is presented in Eq 13 below.

$$W_{\text{Engine}}\text{lbm} = \frac{T_{\text{TO}}\text{lbf}}{5.5} \tag{13}$$

Raymer [2]: Raymer's seventh edition of *Aircraft Design, A Conceptual Approach* includes an estimation method for UEW (source Eq. 10.4), and reproduced below in Eq. 14 for convenience.

$$W_{\text{Engine}}\text{kg} = 14.7 \times T_{\text{TO}}^{1.1}\text{kN} \times e^{-0.045 \times BPR} \tag{14}$$

Jenkinson [12]: Jenkinson's *Civil Jet Aircraft Design* (page 139) gives a relation for "bare" engine mass (equivalent to the UEW nomenclature presented in this work), which is shown in Eq.15 below.

$$W_{\text{Engine}}\text{kg} = 8.7\ \ 1.14 \times BPR \times T_{\text{TO}}\text{kN} \tag{15}$$

Svoboda [35]: *Turbofan engine database as a preliminary design tool*, published in 2000, uses regression analysis to derive a series of engine design equations, including an engine weight estimation, shown in Eq. 16:

$$W_{\text{Engine}}\text{lbm} = 250\ \ 0.175 \times T_{\text{TO}}\text{lbf} \tag{16}$$

Similar to OEW, predicted versus actual plots for UEW regression methods are illustrated in Fig. 6, accompanied by the moments of each method's error distribution, shown in Table 4. The results show that FAST outperforms most methods with respect to mean, median, and standard deviation. FAST's error distribution is both more skewed and has thicker tails than the Raymer and Jenkinson methods. These methods, however, tend to underpredict engine weight, marginalizing their low skewness and kurtosis.

The tendency of FAST and the multivariate linear fit to overpredict while other methods underpredict could be a result of engine over-sizing in recent years. Arnson et al. [9] explain how turbofan engines have begun to be oversized due to bypass ratio increases, driven by the fuel efficiency benefits this brings. These increases have primarily been achieved by engine up-scaling rather than core miniaturization. This results in an engine which, while efficient, may be oversized for its application. This trend could contribute to differences between FAST and historical methods, which may have trained their models using data which does not include engines produced in the recent past.

Table 4: Error metric summary for FAST and reference regressions on engine weight.

Method	Sample Size	Mean (%)	Median (%)	Std. Dev. (%)	Skewness	Kurtosis
FAST	330	+1.010	+1.592	10.74	-0.1357	3.246
Multivariate Linear Fit [33]	330	+5.096	+1.637	31.64	+6.392	51.62
FLOPS[3]	350	-7.662	-8.106	14.72	+0.3769	3.404
Raymer [2]	318	-5.117	-7.759	18.32	-0.0859	2.619
Jenkinson[12]	318	-26.23	-25.13	18.71	-0.0547	2.667
Svoboda [35]	350	-4.799	-6.236	19.38	+2.768	15.56

3.3 Electrified Aircraft Weight Validation

Comparisons to conventional aircraft data provide confidence that the proposed regression methodology is functioning as intended. Additionally, the FAST tool is also intended for use with electrified aircraft concepts, and for this purpose, the regressions should be validated against contemporary designs. This study considers the following aircraft concepts:

NASA's SUSAN: The **SU**bsonic **S**ingle **A**ft e**N**gine (SUSAN) is a hybrid, partially turboelectric aircraft concept, which can be compared to a Boeing 737-8 [31]. It has a 2,500 nautical mile design range, and carries 180 passengers.

Boeing's SUGAR Volt: The **S**ubsonic **U**ltra **G**reen **A**ircraft **R**esearch Volt is a parallel hybrid electric aircraft concept, which has several variants. It has a 3,500 nautical mile design range and carries 154 passengers [36, 37]. The source gives sufficient information to use FAST's regression for two variants, the "low power balanced," and the "core shutdown" which rely on a conservative and more aggressive level of electrification, respectively.

MIT's AEA-800: The **A**ll **E**lectric **A**ircraft 800 (AEA-800) is a fully electric aircraft concept which can be compared to the Airbus A320neo. It has a 500 nautical mile design range, and carries 180 passengers [38].

Elysian's E9X: The E9X, a fully electric aircraft concept, has a design range of 431 nautical miles (800 kilometers), and carries 90 passengers [39].

Each concept's ASO weight is predicted according to the methodology described in Sec. 2.5.1, which requires burden weight, MTOW, range, and SLS thrust. The SUSAN and SUGAR concepts are described comprehensively in references [31] and [37] respectively. Values for the AEA-800 and E9X are inferred from figures, which may introduce small inconsistencies between the sources' true values and the ones listed here. The inputs to the FAST regression are presented in Table 5, while the ASO weight predictions, the source ASO weights, and the differences are reported in Table 6.

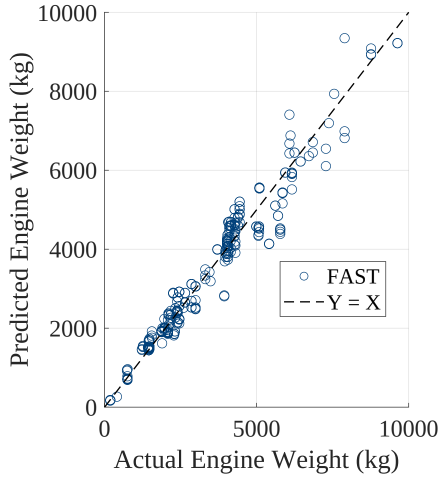

(a) FAST's GPR.

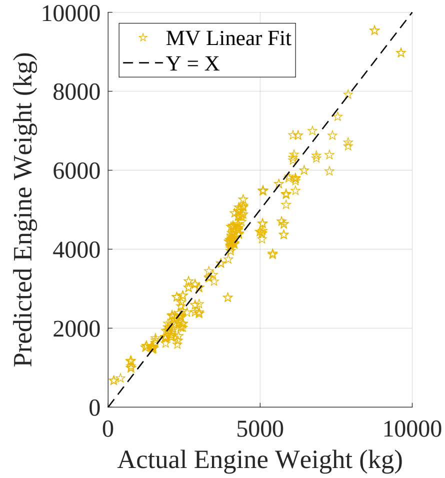

(b) Multivariate linear fit [33].

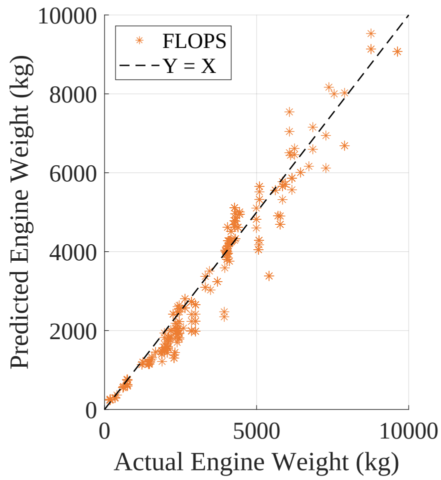

(c) FLOPS equation [3].

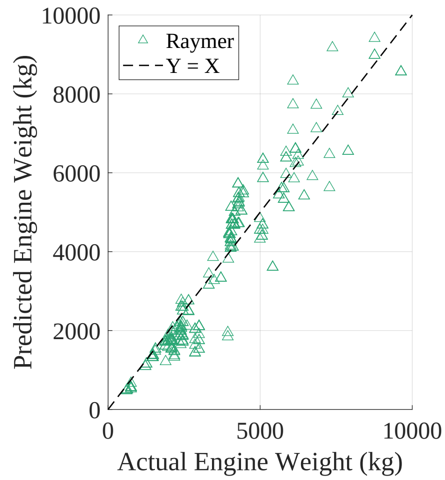

(d) Raymer regression [2].

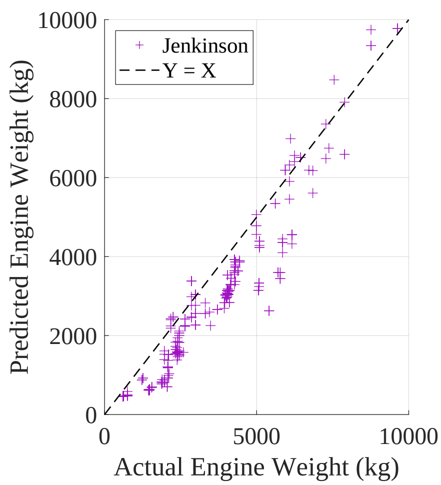

(e) Jenkinson regression [12].

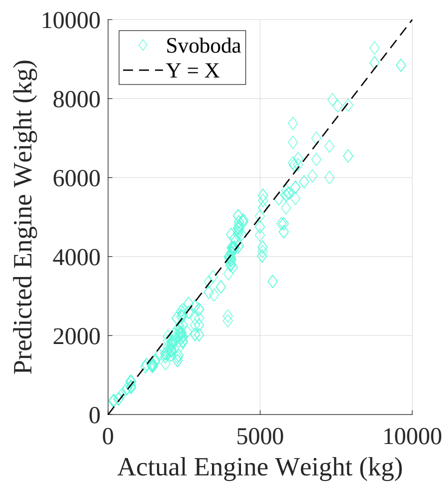

(f) Svoboda regression [35].

Figure 6: Engine weight predictions vs actual for several estimation methods.

The AEA-800 and E9X include estimates for structural weight reductions due to use of composite materials. The AEA-800 assumes that “largely due to the advanced technology assumptions, the fuselage, wing, and tails are over 15%, 50%, and 60% lighter respectively” [38] for the conventional reference configurations. “These benefits are also observed for the AEA-800” [38] (Sec. 5.4). The E9X assumes that both “a 10% and 20% reduction in fuselage mass and tail mass are assumed due to the use of composites, respectively," as well as “when estimating the wing weight ... the max zero fuel mass is taken as the total aircraft mass excluding battery mass. This is based on the assumption that the battery contributes to a wing bending relief, similarly to fuel. A more detailed structural analysis of the wing is required to verify this approach. Additionally, the 10% reduction in wing mass due to bending relief of the engines... is applied, and a 4% wing mass penalty is assumed for the folding wingtips” [39] (Appendix A.B). Both the AEA-800 and E9X assume weight reduction factors larger than those recommended in reference [2].

The composite assumptions are a plausible contributor to the large ASO weight prediction differences for the AEA-800 and E9X (+52% and +72% respectively). The FAST Aerobase consists of historical aircraft, where composite use was not common until recently; consequently FAST has limited ability to extrapolate for future technology assumptions. References [2] and [40], the latter of which is used by reference [39] in the E9X weight buildup, recommend using “fudge factors[5]” for estimating structural weight reductions through the use of composites, which would need to be applied after the ASO weight prediction in FAST. The FAST regression may over-predict the E9X weight in part due to the assumed use of composites, but the larger difference could also arise from dissimilarities in how the two methods consider battery weight when sizing structural weights. FAST assigns battery weight (and fuel weight) as a component in the burden group, while the methodology presented in de Vries and Wolleswinkel [39] assumes that the battery’s weight contribution will relieve bending in the wing, as fuel would in a kerosene-based aircraft.

The differences between FAST and the weights reported in the SUSAN, SUGAR balanced, and SUGAR core-shutdown concepts (-1.2%, -7.0%, and +6.6% respectively) are comparatively much smaller than those from the AEA-800 and E9X. The SUGAR aircraft were developed beyond the conceptual design stage, and while they did assume composite use in their weight buildups, the purpose was to show that advanced composites would be required to support the high aspect ratio seen in the concept without additional weight penalties [36]. Assuming lighter materials at the same time as more demanding structural loads could introduce counteracting influences on OEW, allowing FAST to more accurately predict OEW despite the unconventional architecture and material use. Without developed concepts or prototypes for such aircraft, it is challenging to verify the methodology with as much confidence as the conventional aircraft database can provide. The methodology presented in this work is most robust for aircraft using unconventional propulsion and energy system architectures, and less robust for those using unconventional aerodynamic or structural systems. However, as new data is produced in coming years, adding inputs to the OEW (ASO) regression, such as aspect ratio or composite usage, will keep the regressions calibrated to physical trends, improving their robustness for such aircraft.

Validating FAST’s ASO weight regression with electrified propulsion architectures is less straightforward than that for conventional aircraft, as “turbofan-class” electrified aircraft have yet to be produced. With only conceptual and preliminary designs, often predicated on technological advancements, as points of comparison,

[5]Raymer [2] includes these values in Table 15.4, where the range of weight reductions are 10-15%, 12-17%, 5-10%, 0-5%, and 10-15% for wings, tails, fuselage/nacelle, landing gear, and air induction systems respectively. As previously mentioned, the AEA-800 and the E9X assume fudge factors outside of these ranges.

results should be considered provisional. However, maximum differences of approximately 7% (for aircraft without fudge-factor weight reductions) appear promising for early-stage conceptual design of electrified aircraft.

Table 5: FAST ASO weight regression inputs for electrified concepts.

Aircraft	Burden (kg)	MTOW (kg)	Range (km)	Thrust (kN)
SUSAN	33,720	89,810	4,630	262.5
SUGAR Volt: Balanced	31,970	68,040	6,482	154.8
SUGAR Volt: Core-Shutdown	47,890	86,180	6,482	161.5
AEA-800	65,730	109,500	926	322.3
E9X	50,050	76,000	800	223.7

Table 6: FAST ASO weight regression predictions for electrified concepts.

Aircraft	Predicted ASO (kg)	True ASO (kg)	Difference (%)	Notes
SUSAN	39,310	39,800	-1.231	–
SUGAR Volt: Balanced	33,530	36,070	-7.042	–
SUGAR Volt: Core-Shutdown	40,820	38,290	+6.607	–
AEA-800	65,590	43,250	+51.65	Values inferred from figures, Composites
E9X	44,560	25,950	+71.71	Values inferred from figures, Composites, Battery

4 Conclusions and Future Work

This work presents a novel approach to predicting key aircraft design parameters by leveraging Gaussian Process Regressions applied to a large, open-source historical database of aircraft and engine specifications. The "Aerobase," comprising over 400 aircraft and 200 engines, is made freely available through the Future Aircraft Sizing Tool (FAST), offering the aviation design community an accessible, valuable resource for conceptual aircraft design.

The research demonstrates that the GPR models, integrated with FAST's database, provide improvements over reference regressions from canonical literature. These models incorporate a broader range of input parameters than conventionally used at the early conceptual design stage. The results show that FAST's regression models not only reduce prediction errors, but also deliver tighter error distributions and smaller standard deviations compared to existing methods, providing aircraft designers with more reliable estimates early in the conceptual design phase. In Sec. 2.5.1, "burden" and "ASO" weight groupings were introduced, which allowed the Aerobase data to predict structural weights for electrified aircraft concepts. These predictions were shown to be within approximately -7% and +6% of the values reported in primary sources in the absence of advanced technological assumptions.

A key contribution of this work is the flexibility of FAST's non-parametric GPR framework, which allows for dynamic regression generation without predefined input constraints. This capability enables aircraft designers to explore a wide range of design configurations with precision, demonstrating the tool's adaptability across different aircraft and engine architectures. Additionally, this paper introduced a methodology to adapt conventional data for use in hybrid-electric and fully electric aircraft designs. By modifying conventional regressions, this approach ensures that historical data can still be leveraged to predict parameters for novel propulsion systems, enabling the design of sustainable, electrified aircraft concepts.

The development of the Aerobase and regression tools is part of an ongoing effort to continually refine and expand the FAST platform. Future work will focus on incorporating additional aircraft and engine data as they become available, and further improving the accuracy of the regression models as new user-contributed data enhances the database. Moreover, the flexibility of FAST allows it to be adapted for various aircraft configurations and propulsion technologies, including the integration of alternative fuels and novel electrified propulsion architectures. As the tool evolves, it is expected to provide more robust capabilities for the conceptual design of next-generation aircraft.

Appendix

4.1 FAST Regression Tutorial

Figure 7 shows a very simple example of a user specified regression call in FAST, formatted using Reference [41]. This example illustrates a prediction for OEW (kg) based on a known MTOW (kg) and aircraft range (m). FAST aircraft are stored in a data structure, so each parameter on lines 2-4 are the paths through the structure to the location of the desired parameter. FAST will always treat the last entry of the Input/Output cell array as the output. Input order does not matter so long as the Target array matches the variable order, shown in line 8. Line 8 declares the input parameter values that are known by a user, and as mentioned the order must match that of the Input. In this example, since MTOW is declared as the first input while Range is declared as the second, and MTOW = 50,000 kg and Range = 3,000 km,

`Target = [5e4 , 3000e3]` is correct, while

`Target = [3000e3 , 5e4]` would be *incorrect*.

With the variables and their values declared, the next step is to load the FAST database, shown in line 11. The database stores 6 structures: turbofan aircraft, turbofan engines, turboprop aircraft, turboprop engines, fan reference, and propeller reference. The latter two mimic the structure of the others such that the format of the data structures can be explored and the variable paths (lines 2-4 in the example) can be known. The reference structures contain units in place of the numerical values that are stored in the database, so targets (line 8) can be entered in the correct units. The engine database contains all engines which are used by the aircraft in the aircraft database. The same information for each engine is stored within `Aircraft.Specs.Propulsion.Engine`. The engine databases are included separately because a user may wish to run a regression on an engine parameter (such as OPR as a function of EIS), and using the aircraft database would be biased towards engines which are used more frequently among different types of aircraft.

With all required input variables initialized, the regression can be run, as shown on line 15. Outputs of the regression function are the mean and variance of the normal distribution, which the output parameter is assumed to take. In FAST, the mean is always taken as the value the regression predicts while the variance is used as a measure of confidence the regression has in its prediction. Two additional inputs, weights and prior, can be included after a target in the regression call. Both are options to manually tune the regression and are not used in FAST. Weights adds a scaling factor to each input variable when summed in the covariance matrix as described in Sec. 2.4. In this case, replacing line 15 with

```
[Mu, Sigma2] = RegressionPkg.NLGPR(TurbofanAC, InputOutput, Target, [2 , 1]);
```

would tell the regression that MTOW is twice as important as range when predicting OEW. Finally, a prior can be set which overwrites the FAST default of the database average. If a user believed the OEW in this example should be 20,000 kg, but wanted the database to modify that guess using historical data, they would replace line 15 with

```
[Mu, Sigma2] = RegressionPkg.NLGPR(TurbofanAC, InputOutput, Target, [1 , 1], 20e3);
```

Finally, the regression is capable of making multiple predictions in a single call. The only change to the inputs would be in Target and Prior (if Prior was set). Both of the variables would get extended to be multiple *rows* larger. For example, if a user wanted to additionally predict the OEW at an MTOW of 40,000 kg and 2,500 km, line 8 would be replaced with

```
Target = [50e3, 3000e3; 40e3, 2500e3];
```

whereupon the outputs would be 2×1 sized vectors as opposed to scalars that the example yields.

```
1  % Set Input/Output Variables (Cell array composed of string vectors)
2  InputOutput = {["Specs", "Weight",      "MTOW" ],...
3                 ["Specs", "Performance", "Range"],...
4                 ["Specs", "Weight",      "OEW"  ]};
5
6  % Set known input space values (double vector)
7  %          MTOW       Range
8  Target = [50e3,      3000e3];
9
10 % Load database (TurbofanAC variable will be loaded)
11 load('+DatabasePkg/IDEAS_DB.mat')
12
13 % Call the regression, predicting a mean and variance for the distribution
14 % of OEWs at the given Target vector
15 [Mu, Sigma2] = RegressionPkg.NLGPR(TurbofanAC, InputOutput, Target);
```

Figure 7: Example single regression call in FAST.

4.2 FLOPS Verification

FLOPS is the most involved of the OEW prediction methodologies. For reference, this appendix includes assumptions which are made by the FLOPS methodologies, as well as comparisons of the adapted version to FLOPS 8.11 for two notional aircraft. These vehicles are described by a set of input parameters which are not necessarily representative of any true aircraft. They are intended to show agreement between the two methods using prototypical values for a single aisle and a twin aisle aircraft.

FLOPS expects several parameters which may not be listed in the database. Table 7 provides a comprehensive list of all parameters which are input into both the MATLAB adapted version, as well as the the source of the value, and the value used if it was an assumed parameter. FLOPS uses default values for undeclared parameters, and values are expected in Imperial units. Tables 8 and 9 show the input space and resultant comparison between the MATLAB adapted version and the original FLOPS 8.11, respectively, for the notional single aisle aircraft. Tables 10 and 11 are repeated for the notional twin aisle case. Both the verification aircraft show group weights within 5% agreement and OEW within 1%, which this work considers sufficient to declare the MATLAB adaptation a good recreation of FLOPS 8.11.

FLOPS was adapted in order to avoid the necessity for a wrapper (FLOPS input file builder). As the authors are less familiar with the syntax, a recreation in a familiar environment was more convenient to implement than a wrapper.

Table 7: Input Space for FLOPS

Parameter	FLOPS Name	Source	Value
% SW used for Control Surfaces	FLAPR	[3]	0.333
# Fuselage Engines	NEF	Aerobase	–
# Wing Engines	NEW	Aerobase	–
Aeroelastic Tailoring Flag	FAERT	Assumed	False
Business Class Passengers	NPB	Assumed/Aerobase	NPB = # Pax
Cabin Length	XLP	[3]	Eq. 208
Composite Use Flag	FCOMP	Assumed	False
Cruise Mach Number	VCMN	Assumed	0.8
Design Range	DESRNG	Aerobase	–
Dihedral	DIH	Assumed	3°
Fuel Weight	FULWMX	Aerobase	–
Fuselage Depth	DF	Assumed	DF = Height/3
Fuselage Length	XL	Aerobase	–
Fuselage Width	WF	Assumed	WF = Height/3
Horizontal Tail Volume Coeff.	HTVC	[3]	Eqs.173-178
MTOW	GW	Aerobase	–
Nacelle Diameter	DNAC	Assumed/Aerobase	DNAC = Fan Diameter
Nacelle Length	XNAC	Assumed/Aerobase	XNAC = Engine Length
Number of Fuel Tanks	NTANK	Assumed	3
Paint Density	WPAINT	Assumed	0.03 lb / ft^2
Quarter Chord Sweep	SWEEP	Aerobase	–
Taper Ratio	TR	Aerobase	–
Thrust per Engine	THRSO	Aerobase	–
Thrust Reverser Flag	WTHR	Assumed	True
Variable Sweep Flag	VARSWP	Assumed	False
Vertical Tail Volume Coeff.	VTVC	[3]	Eqs.173-178
Weighted Average Thickness-to-Chord	TCA	[3]	0.15
Wing Area	SW	Aerobase	–
Wing Strut Flag	FSTRT	Assumed	False
Wingspan	SPAN	Aerobase	–

Funding Sources

This work is sponsored by the NASA Aeronautics Research Mission Directorate and Electrified Powertrain Flight Demonstration project, "Development of a Parametrically Driven Electrified Aircraft Design and Optimization Tool", WO-0238. Karin Bozak serves as NASA's Technical Monitor.

Table 8: Input Values for the notional single aisle verification.

Parameter	FLOPS Name	Value
# Fuselage Engines	NEF	0
# Wing Engines	NEW	2
Business Class Passengers	NPB	150
Cabin Length	XLP	160 ft
Design Range	DESRNG	2,500 nmi
Fuselage Depth	DF	20 ft
Fuselage Length	XL	200 ft
Fuselage Width	WF	20 ft
Fuel Weight	FULWMX	50,000 lbm
Horizontal Tail Volume Coeff.	HTVC	3.44
MTOW	GW	250,000 lbm
Nacelle Diameter	DNAC	8 ft
Nacelle Length	XNAC	12 ft
Quarter-Chord Sweep	SWEEP	30°
Taper Ratio	TR	0.20
Thrust per Engine	THRSO	25,000 lbf
Vertical Tail Volume Coeff.	VTVC	0.136
Wing Area	SW	3,500 ft^2
Wingspan	SPAN	200 ft

Table 9: Comparison of FLOPS methodologies for the notional single aisle aircraft.

Weight Group	Adapted (lbm)	FLOPS 8.11 (lbm)	Difference (%)
Structural	1.123e+05	1.122e+05	+0.0169
Propulsion	12020	12150	-1.087
Systems and Equipment	44920	44580	+0.7748
Operations	6254	6459	-3.181
OEW (total)	1.755e+05	1.754e+05	+0.01582

Table 10: Input Values for the notional twin aisle verification.

Parameter	FLOPS Name	Value
# Fuselage Engines	NEF	1
# Wing Engines	NEW	2
Business Class Passengers	NPB	250
Cabin Length	XLP	260 ft
Design Range	DESRNG	4,000 nmi
Fuselage Depth	DF	30 ft
Fuselage Length	XL	300 ft
Fuselage Width	WF	30 ft
Fuel Weight	FULWMX	75,000 lbm
Horizontal Tail Volume Coeff.	HTVC	5.86
MTOW	GW	550,000 lbm
Nacelle Diameter	DNAC	12 ft
Nacelle Length	XNAC	16 ft
Quarter-Chord Sweep	SWEEP	27°
Taper Ratio	TR	0.2
Thrust per Engine	THRSO	40,000 lbf
Vertical Tail Volume Coeff.	VTVC	0.21
Wing Area	SW	5,500 ft^2
Wingspan	SPAN	250 ft

Table 11: Comparison of FLOPS methodologies for the notional twin aisle aircraft.

Weight Group	Adapted (lbm)	FLOPS 8.11 (lbm)	Difference (%)
Structural	3.03e+05	3.023e+05	+0.2169
Propulsion	29000	29440	-1.49
Systems and Equipment	86890	90620	-4.114
Operations	11310	11820	-4.343
OEW (total)	4.302e+05	4.342e+05	-0.9269

Acknowledgments

The authors would like to thank Ralph Jansen, Amy Chicatelli, Andrew Meade, Karin Bozak, Noah Listgarten, Dennis Rohn, and Gaudy Bezos-O'Connor from NASA's Electrified Powertrain Flight Demonstration project for supporting this work and providing valuable technical input and feedback throughout the duration of the project. The authors also thank Huseyin Acar, Benjamin Bates, Nawa Khailany, Paul Mokotoff, Janki Patel, Joaquin Rey, Jayda Shine, and Michael Tsai for their contributions to either FAST or the historical aircraft database that the regressions and projections rely on. Finally the authors thank Shugo Kaneko for his discussion regarding the FLOPS software.

Generative AI, specifically ChatGPT-5, was used in the proofing of this manuscript. No AI generated images or code was used for this research.

References

[1] Jan Roskam. *Airplane Design: Preliminary Sizing of Airplanes*. DAR Corporation, 1985.

[2] Daniel P Raymer. *Aircraft Design, a Conceptual Approach*. American Institute of Aeronautics & Astronautics, Reston, VA, 7th edition, July 2024.

[3] Douglas P. Wells, Bryce L. Horvath, and Linwood A. McCullers. The flight optimization system weights estimation method. NASA Technical Memorandum NASA/TM–2017–219627/Vol. I, National Aeronautics and Space Administration, Langley Research Center, Hampton, VA, June 2017. Errata issued April 2018.

[4] Jennifer Gratz, Jason Kirk, Carl Recine, John Jasa, Eliot Aretskin-Hariton, Kenneth Moore, and Kiran Marfatia. Aviary: An open-source multidisciplinary design, analysis, and optimization tool for modeling aircraft with analytic gradients. In *AIAA AVIATION FORUM AND ASCEND 2024*. American Institute of Aeronautics and Astronautics, July 2024.

[5] Gokcin Cinar, Yu Cai, Mayank V. Bendarkar, Andrew I. Burrell, Russell K. Denney, and Dimitri N. Mavris. System analysis and design space exploration of regional aircraft with electrified powertrains. *Journal of Aircraft*, 60(2):382–409, March 2023.

[6] Maxfield G Arnson, Gokcin Cinar, Elias Waddington, Phillip Ansell, Matthew Clarke, Reynard de Vries, Francesco Salucci, Nirmit Prabhakar, Jonathan C Gladin, Mingxuan Shi, Edward Lovelace, and Peter de Bock. System-level energy pack requirements for sustainable commercial aviation. In *AIAA AVIATION FORUM AND ASCEND 2024*, page 3828. AIAA, jul 2024.

[7] IDEAS Laboratory. Future aircraft sizing tool (fast). *Github*, 2024.

[8] Paul Mokotoff, Maxfield Arnson, and Gokcin Cinar. Fast: A future aircraft sizing tool for electrified aircraft design. AIAA SciTech Forum, 2025.

[9] M. Arnson, H. Acar, M. Tsai, and G. Cinar. Historical trends and future projections of key performance parameters in aircraft design. Manuscript under review, Journal of Aircraft, 2025.

[10] Nathan Meier. Jet engine specification database. retrived from `https://www.jet-engine.net/`, Sep 2021.

[11] Lloyd Jenkinson, Paul Simpkin, and Darren Rhodes. Civil jet aircraft design data sets, 2001.

[12] Lloyd Jenkinson, Paul Simpkin, and Darren Rhodes. *Civil Jet Aircraft Design*. AIAA Inc., January 1999.

[13] Janes Information Services. *Janes All the World's Aircraft: In Service*. Janes Group, 2025.

[14] Janes Information Services. All the world's aircraft, 2025.

[15] Eurocontrol. Aircraft performance database. retrived from `https://contentzone.eurocontrol.int/aircraftperformance/default.aspx`, 2023.

[16] Eurocontrol. Base of aircraft data (bada), June 2025.

[17] Élodie Roux. *Turbofan and Turbojet Engines: Database Handbook*. Éditions Élodie Roux, 2007.

[18] Élodie Roux. *Turboshaft, Turboprop and Propfan: Database Handbook*. Éditions Élodie Roux, 2011.

[19] Federal Aviation Administration. Aircraft characteristics data. retrieved from FAA, `https://www.faa.gov/airports/engineering/aircraft_char_database`, 2023.

[20] Federal Aviation Administration. Type certificate data sheet a16we, revision 66, 2020.

[21] Boeing Commercial Airplanes. Using the payload/range and takeoff field length charts in the airplane characteristics for airport planning documents, Feb 2014.

[22] Gas turbine engines. *Aviation Week and Space Technology*, January 2008.

[23] Yashovardhan S. Chati and Hamsa Balakrishnan. Modeling of aircraft takeoff weight using gaussian processes. *Journal of Air Transportation*, 26(2):70–79, April 2018.

[24] Nicolas F. M. Wahler, Daigo Maruyama, and Ali Elham. Credibility-based multidisciplinary design optimization of hybrid and fully electric aircraft. *Journal of Aircraft*, page 1–17, January 2025.

[25] Sooryun Shin, Sanga Lee, Chankyu Son, and Kwanjung Yee. Aircraft design parameter estimation using data-driven machine learning models. In *33rd Congress of the International Council of the Aeronautical Sciences, ICAS 2022*, 33rd Congress of the International Council of the Aeronautical Sciences, ICAS 2022, pages 449–465. International Council of the Aeronautical Sciences, 2022. Publisher Copyright: Copyright © (2022) by International Council of Aeronautical Sciences (ICAS) All rights reserved.; 33rd Congress of the International Council of the Aeronautical Sciences, ICAS 2022 ; Conference date: 04-09-2022 Through 09-09-2022.

[26] Eric Schulz, Maarten Speekenbrink, and Andreas Krause. A tutorial on gaussian process regression: Modelling, exploring, and exploiting functions. *Journal of Mathematical Psychology*, 85:1–16, 2018.

[27] John Shawe-Taylor and Nello Cristianini. *Basic kernels and kernel types*, page 291–326. Cambridge University Press, 2004.

[28] David John Cameron MacKay. Introduction to gaussian processes. 1998.

[29] Carl Edward Rasmussen and Christopher K I Williams. *Gaussian Processes for Machine Learning*. Adaptive Computation and Machine Learning series. MIT Press, London, England, November 2005.

[30] Justin S. Gray, John T. Hwang, Joaquim R. R. A. Martins, Kenneth T. Moore, and Bret A. Naylor. OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 59(4):1075–1104, April 2019.

[31] Timothy Chau and Jared Duensing. Conceptual design of the hybrid-electric subsonic single aft engine (susan) electrofan transport aircraft. In *AIAA SCITECH 2024 Forum*. American Institute of Aeronautics and Astronautics, January 2024.

[32] Lloyd R. Jenkinson and James F. Marchman. *Aircraft Design Projects for Engineering Students*. Butterworth-Heinemann, 2003.

[33] The MathWorks, Inc. *MATLAB and Statistics and Machine Learning Toolbox Release R2022b*. The MathWorks, Inc., Natick, Massachusetts, 2022. `https://www.mathworks.com/products/statistics.html`.

[34] Stephen G. West, Janet F. Finch, and Patrick J. Curran. Structural equation modeling with non-normal and missing data. In Rick H. Hoyle, editor, *Structural Equation Modeling*, chapter 4, page 74. SAGE Publications, Thousand Oaks, CA, May 1995.

[35] Charlie Svoboda. Turbofan engine database as a preliminary design tool. *Aircraft Design*, 3(1):17–31, 2000.

[36] Marty K. Bradley and Christopher K. Droney. Subsonic ultra green aircraft research: Phase ii - volume i - truss braced wing design exploration. Technical Report NASA/CR–2015-218704, NASA, Huntington Beach, California, 2015.

[37] Marty K. Bradley and Christopher K. Droney. Subsonic ultra green aircraft research: Phase ii - volume ii - hybrid electric design exploration. Technical Report NASA/CR–2015-218704, NASA, Huntington Beach, California, 2015.

[38] Albert R. Gnadt, Raymond L. Speth, Jayant S. Sabnis, and Steven R.H. Barrett. Technical and environmental assessment of all-electric 180-passenger commercial aircraft. *Progress in Aerospace Sciences*, 105:1–30, February 2019.

[39] Reynard de Vries, Rob E. Wolleswinkel, Maurice Hoogreef, and Roelof Vos. A new perspective on battery-electric aviation, part ii: Conceptual design of a 90-seater. In *AIAA SCITECH 2024 Forum*. American Institute of Aeronautics and Astronautics, January 2024.

[40] Egbert Torenbeek. *Synthesis of Subsonic Airplane Design*. Delft University Press, Delft, 1982.

[41] Florian Korn. mcode.sty demo. `https://www.overleaf.com/latex/templates/highlighting-matlab-code-in-latex-with-mcode/nhtksndnsmmx.pdf`, 2014.