



## salmon: A Symbolic Linear Regression Package for Python

Alex Boyd 

University of California, Irvine

Dennis L. Sun 

California Polytechnic State University

---

### Abstract

One of the most attractive features of R is its linear modeling capabilities. We describe a Python package, **salmon**, that brings the best of R's linear modeling functionality to Python in a Pythonic way – by providing composable objects for specifying and fitting linear models. This object-oriented design also enables other features that enhance ease-of-use, such as automatic visualizations and intelligent model building.

*Keywords:* linear regression, linear model, visualization, model building, Python.

---

## 1. Introduction

Linear models are ubiquitous in statistics, science, engineering, and machine learning. A linear model assumes that the expected value of a response variable  $y$  is a linear function of explanatory variables,  $x_1, \dots, x_p$ :

$$E[y \mid \mathbf{x}] = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_p x_p \quad (1)$$

for coefficients  $\beta_0, \beta_1, \dots, \beta_p$ . Model (1) is more flexible and general than it may first appear. For example, linear models do not necessarily have to be linear in the original explanatory variables. If we add higher-order polynomial terms to a linear model, then we can also model non-linear effects:

$$E[y \mid \mathbf{x}] = \beta_0 + \beta_1 x_1 + \beta_2 x_1^2 + \beta_3 x_1^3 + \beta_4 x_2. \quad (2)$$

How are linear models used? First, they can be used for *description*. For example, we may want to emphasize trends in a scatterplot by superimposing a best-fit line on the points. Second, they may be used for *prediction*. Once the coefficients  $\beta_0, \dots, \beta_p$  have been estimated, the linear model can be used to predict the value of the response  $y$  for a new observation where only the explanatory variables are known. Finally, linear regression can be used for

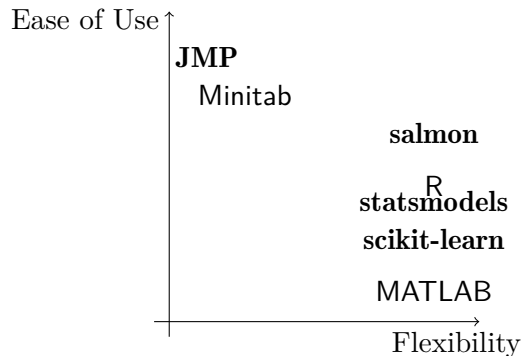


Figure 1: Tradeoffs of different solutions for fitting linear models.

*inference*. The coefficients may encode information about nature, such as the causal effect of one variable on another, in which case we need hypothesis tests and confidence intervals for model parameters.

Because of these different use cases, several solutions for fitting linear models have emerged. Figure 1 illustrates the tradeoffs. At one extreme are point-and-click software packages, like **JMP** (SAS Institute Inc. 2019) and **Minitab** (Minitab Inc. 2017), which fit linear models and provide automatic visualizations. Although these packages have much built-in functionality, they are not easily extensible. For example, complicated data cleaning and wrangling often have to be done outside these software packages.

At the other extreme are programming languages for scientific computing, like **MATLAB** (The MathWorks Inc. 2021) and **Python** (Van Rossum *et al.* 2011). To fit a linear model, users have to manually construct the matrices to be passed to a least-squares solver. To obtain predictions from the fitted model, users must implement the matrix multiplications. Although these environments can be powerful, users have to keep track of low-level details that distract from the modeling.

Libraries within these languages ease the burden somewhat. For example, **scikit-learn** is a Python library for machine learning that provides a consistent API (application programming interface) for specifying, fitting, and predicting using a linear regression model (Pedregosa *et al.* 2011). However, it provides little help with the other two uses of linear regression, description and inference, offering neither visualizations nor uncertainty estimates. Also, the `LinearRegression` model in **scikit-learn** assumes that the explanatory variables have already been transformed into the numerical matrix that will be passed into a least-squares solver, so users must manually transform the variables or else define the transformations as part of the model.

**R** (R Core Team 2023) occupies a medium between the two extremes. On the one hand, it is a full-fledged programming language. On the other, it provides a high-level API for specifying and fitting linear regression models through formulas and the `lm` function. For example, Model 2 could be specified and fit in R as

```
R> lm(y ~ poly(x1, 3) + x2, data)
```

Although R provides automatic diagnostic plots, it offers limited visualizations of the fitted model, in comparison with point-and-click software packages such as **JMP** and **Minitab**.

**patsy** (Smith *et al.* 2022) and **statsmodels** (Seabold and Perktold 2010) are Python libraries that port R-style modeling to Python. Like R, **statsmodels** provides some automatic diagnostic plots but not visualizations of the fitted model. However, R formula syntax is not interpretable by Python, so formulas have to be specified as strings:

```
>>> smf.ols("y ~ poly(x1, 3) + x2", data)
```

This creates problems when the column names do not meet the language’s rules for variable names. For example, columns with names such as "weight.in.kg" or "person’s height" would need to be wrapped in **patsy**’s “quote” object, **Q**:

```
>>> smf.ols("y ~ Q('weight.in.kg') + Q('person\'s height')", data)
```

Because the **Q** constructor takes a string as input, we have strings within strings – hence the need for two different kinds of string delimiters. If the variable name itself includes an apostrophe or a quotation mark, then that character needs to be escaped. An object-oriented approach to model specification, described in this paper, avoids these complications, allowing any object that is a valid column name in a **pandas DataFrame** to be used directly in the model specification, without requiring a workaround for problematic variable names. Furthermore, having objects associated with each variable in a model makes it easy to specify customization options for each variable (e.g., the baseline level for a categorical variable).

In this paper, we describe **salmon**, a Python package for linear regression that offers an object-oriented interface for specifying linear models. The two main contributions of **salmon** are:

1. Providing an R-like (but Pythonic) API for specifying models.
2. Producing appropriate visualizations of the models, bridging the gap with point-and-click packages, like **JMP** and **Minitab**.

Its philosophy and design is similar to other statistical packages in Python, such as **Symbulate** (Ross and Sun 2019). Throughout, we provide comparisons of model specification in **salmon** with the similar formula syntax of R.

The easiest way to obtain **salmon** is to install it via **pip**:

```
pip install salmon-lm
```

but it can also be installed from source at <http://github.com/ajboyd2/salmon>.

## 2. Model building and fitting

We introduce the design and syntax of **salmon** by way of a case study. We assume that all **salmon** objects and methods have already been imported into the global namespace, as follows:

```
>>> from salmon import Q, C, Log, Poly, LinearModel
```

	Neighborhood	Price(\$)	Style	Sq. Ft.	Fire?
0	SawyerW	162000	2 Story	1400	No
1	CollgCr	195000	2 Story	1660	No
2	Crawfor	164000	Other	1646	Yes
3	NridgHt	417500	1 Story	2464	Yes
4	SawyerW	186800	1 Story	1400	No

Table 1: First five rows of the Ames Housing dataset, with only the relevant columns for this paper shown.

We will use a sampled version of the Ames housing data set (De Cock 2011), which can be found within the repository for the package. The first five rows of this sampled data set are shown in Table 1.

We will start with the simplest possible model, which assumes the sale price is a linear function of just the square footage:

$$E[\text{Price}(\$) \mid \mathbf{x}] = \beta_0 + \beta_1(\text{Sq. Ft.}). \quad (3)$$

This simple linear regression model can be specified in **salmon** as follows:

```
>>> x = Q("Sq. Ft.")
>>> y = Q("Price($)")
>>> simple_model = LinearModel(x, y)
>>> print(simple_model)
```

```
Price($) ~ 1 + Sq. Ft.
```

Notice that a quantitative variable is specified by creating a **Quantitative** object, or **Q** for short, with the name of the column in the **DataFrame**. Alternatively, we could have created a generic variable using **V** and let **salmon** infer the type. Either way, these variable objects become the explanatory (**x**) and response (**y**) components of a ‘**LinearModel**’ object.

An intercept is added by default, as evidenced by the constant term **1** in the printout. To specify a model without an intercept, we could either insert **- 1** into the expression (which mirrors the formula syntax of **R**) or specify `intercept = False`:


```
>>> no_intercept_model = LinearModel(x - 1, y)
>>> no_intercept_model = LinearModel(x, y, intercept = False)
>>> print(no_intercept_model)
```

```
Price($) ~ Sq. Ft.
```

To fit the `simple_model` above to data, we call the `.fit()` method and pass in a **pandas DataFrame** containing those variables:

```
>>> simple_model.fit(data)
```

	Coefficient	SE	$t$	$p$	2.5%	97.5%
Sq. Ft.	118.5	2.068	57.33	0.000	114.5	122.6
Intercept	3614	3254	1.111	0.2668	-2766	9995

Table 2: Coefficients and inferences for Model 3, `simple_model`.Figure 2: Basic visualization of the fitted Model 3, `simple_model`.

Notice that `.fit()` returns the standard regression output – containing the coefficients, their standard errors, the  $t$  statistic and  $p$  value for testing  $\beta_j = 0$ , and associated 95% confidence intervals – stored in a `DataFrame` for easy display and access. From the regression output, Sq. Ft. appears to have substantial explanatory power, but to be sure, we should visualize the model. A model can be visualized using the `.plot()` method; **salmon** will automatically choose an appropriate visualization:

```
>>> simple_model.plot()
```

Notice that the line of code above produces a scatterplot of the data in black, with the fitted regression line superimposed in red (Figure 2). To produce a similar plot from a model that was fit in R, **scikit-learn**, or **statsmodels**, we would have had to first create a grid of “test”  $x$  values, use the fitted model to predict the value of the response at each of those  $x$  values, and then plot these predictions as a line using some plotting library.

Confidence and prediction intervals can be added to any model visualization, by passing the desired error rate  $\alpha$  (so that the confidence level is  $1 - \alpha$ ) to arguments `confidence_band` or `prediction_band` in `.plot()`:

```
>>> alpha_val = 0.05
>>> simple_model.plot(confidence_band = alpha_val)
```

In Figure 3, the points appear to “fan out” from the fitted model as square footage increases. This suggests that the assumption of constant variance (homoskedasticity) may be violated.




Figure 3: Visualization of the fitted Model 3, `simple_model`, with confidence band.

One fix is to transform the response variable. In other words, we can instead fit the model:

$$E[\log(\text{Price}(\$)) \mid \mathbf{x}] = \beta_0 + \beta_1(\text{Sq. Ft.}) \quad (4)$$

which can be accomplished by literally transforming the response:

```
>>> simple_log_model = LinearModel(Q("Sq. Ft."), Log(Q("Price($)")))
>>> print(simple_log_model)
```

```
log(Price($)) ~ 1 + Sq. Ft.
```

Compare with R, where the same model would be specified as

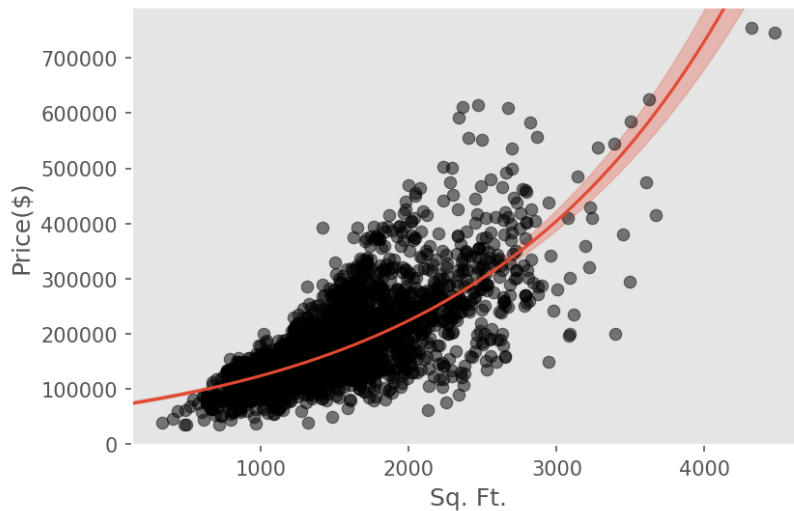
```
R> log(`Price($)` ) ~ `Sq. Ft.`
```

Note that the backticks are only necessary due to the variable names containing special characters (e.g., spaces and dollar sign).

**salmon** supports a number of different transformations by default, such as:

- Natural Logarithm: `Log(X)`.
- Logarithm of Base 10: `Log10(X)`.
- Sine: `Sin(X)`.
- Cosine: `Cos(X)`.
- Exponential: `Exp(X)`.
- Standardization: `Std(X)`, `Standardize(X)`.
- Centering: `Cen(X)`, `Center(X)`.
- Identity: `Identity(X)`.

	Coefficient	SE	$t$	$p$	2.5%	97.5%
Sq. Ft.	5.906e-4	1.055e-5	55.96	0.000	5.699e-4	6.113e-4
Intercept	11.14	0.0166	670.6	0.000	11.11	11.17

Table 3: Coefficients and inferences for the fitted Model 4, `simple_log_model`.Figure 4: Visualization of the fitted Model 4, `simple_log_model`, with confidence band. By default, **salmon** plots the original variables, rather than the transformed variables.

Now we can fit and visualize this new model:

```
>>> simple_log_model.fit(data)
>>> simple_log_model.plot(confidence_band = alpha_val)
```

By default, **salmon** plots the variables in the original (untransformed) space, which is usually desired. After all, we are interested in Price(\$), not  $\log(\text{Price}(\$))$ . However, the log-space view can be helpful for checking whether the linear regression assumptions are met. To produce a plot in the transformed space, specify `transformed_y_space = True`.

```
>>> simple_log_model.plot(confidence_band = alpha_val,
...     transformed_y_space = True)
```

From Figure 5, it is clear that the fanning is indeed reduced in log-space – as we had hoped! However, the model may be further improved by adding a quadratic term to capture the parabolic relationship. We create a new model for this:

$$E[\log(\text{Price}(\$)) \mid \mathbf{x}] = \beta_0 + \beta_1(\text{Sq. Ft.}) + \beta_2(\text{Sq. Ft.})^2 \quad (5)$$

A polynomial model, like Model 5, can be specified in **salmon** using the `Poly()` class, much as one uses the `poly()` function in R:

```
>>> poly_model = LinearModel(Poly(Q("Sq. Ft."), 2), Log(Q("Price($)")))
>>> print(poly_model)
```

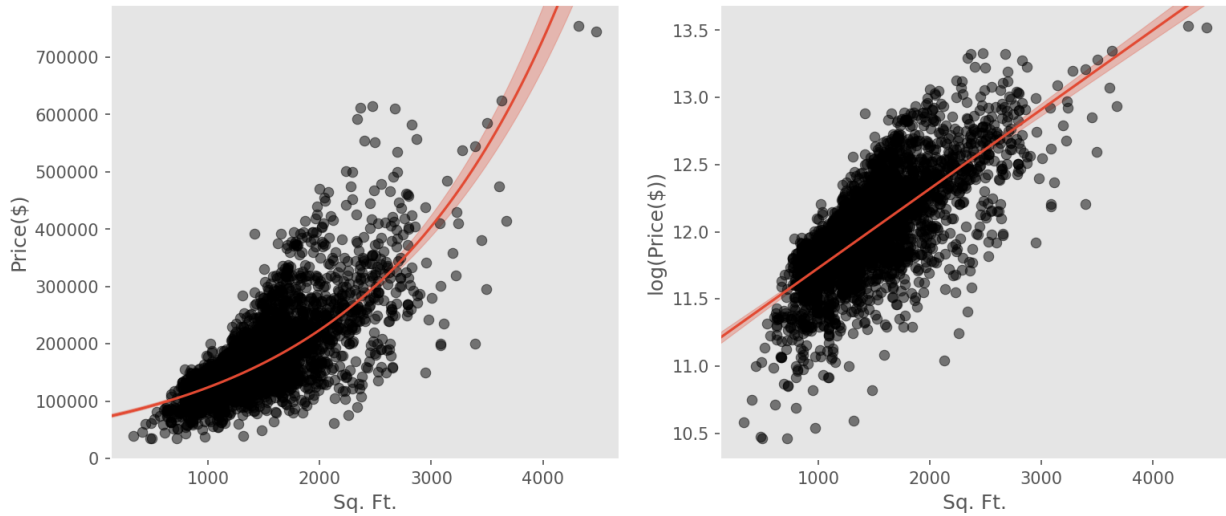


Figure 5: Two visualizations of Model 4, `simple_log_model`. The plot on the left shows the original response (Price(\$)), while the plot on the right shows the transformed response ( $\log(\text{Price}(\$))$ ).

	Coefficient	SE	$t$	$p$	2.5%	97.5%
Sq. Ft.	9.989e-4	4.603e-5	21.70	5.9e-97	9.087e-4	0.0011
(Sq. Ft.) <sup>2</sup>	-1.180e-7	1.296e-8	-9.107	1.5e-19	-1.434e-7	-9.261e-8
Intercept	10.82	0.0386	280.3	0.000	10.75	10.90

Table 4: Coefficients and inferences for Model 5, `poly_model`.

```
log(Price($)) ~ 1 + Sq. Ft. + (Sq. Ft.)^2
```

```
>>> poly_model.fit(data)
```

As can be seen, in this example `Poly(Q("Sq. Ft."), 2)` is equivalent to `Q("Sq. Ft.") + Q("Sq. Ft.") ** 2`.

```
>>> poly_model.plot(confidence_band = alpha_val, transformed_y_space = True)
```

The  $p$  value for the quadratic term (Sq. Ft.)<sup>2</sup> is practically zero, which suggests that the fit is substantially improved by adding the quadratic term.

So far, the models have only used square footage. It is worth investigating whether the fit can be improved by adding the two categorical variables in the data set – the presence of a fireplace and the overall style of house. The former is binary, with a value of “Yes” indicating that the house has a fireplace. The latter can take three different values: “1 Story”, “2 Story” and “Other”.

$$E[\log(\text{Price}(\$)) \mid \mathbf{x}] = \beta_0 + \beta_1 I(\text{Style} = \text{“2 Story”}) + \beta_2 I(\text{Style} = \text{“Other”}) + \beta_3 I(\text{Fire?} = \text{“Yes”}) \quad (6)$$

where  $I(\cdot)$  denotes the indicator function.






Figure 6: Two visualizations of the fitted polynomial Model 5, `poly_model`. The plot on the left shows the original response (Sale Price), while the plot on the right shows the transformed response ( $\log(\text{Price}(\$))$ ).

	Coefficient	SE	$t$	$p$	2.5%	97.5%
Style{2 Story}	0.1058	0.0143	7.374	2.1e-13	0.0777	0.1339
Style{Other}	-0.1500	0.0164	-9.156	9.8e-20	-0.1821	-0.1179
Fire?{Yes}	0.3966	0.0124	31.87	1e-191	0.3722	0.4210
Intercept	11.82	0.0105	1124	0.000	11.80	11.84

Table 5: Coefficients and inferences for Model 6, `simple_cat_model`.

```
>>> simple_cat_model = LinearModel(C("Style") + C("Fire?"),
...   Log(Q("Price($))))
>>> print(simple_cat_model)
```

```
Log(Price($)) ~ 1 + Style + Fire?
```

```
>>> simple_cat_model.fit(data)
```

The same model would be specified in R as

```
R> log(`Price($)` ) ~ `Style` + `Fire?`.
```

The backquotes are necessary because of the non-standard column names. Also, the code above assumes that `Style` and `Fire?` are unambiguously categorical. If they could be confused for quantitative variables, they would have to be explicitly cast to factors:

```
R> log(`Price($)` ) ~ as.factor(`Style`) + as.factor(`Fire?`)
```

This is an example where `salmon`'s insistence on explicit variable types using `C` and `Q` saves typing in the long run.




Figure 7: Two visualizations of Model 6, `simple_cat_model`, which contains two categorical predictors. The plot on the left shows the original response (`Price($)`), while the plot on the right shows the transformed response (`log(Price($))`).

```
>>> simple_cat_model.plot(confidence_band = alpha_val,
...     transformed_y_space = True)
```

To specify that these variables are categorical, we create a `Categorical` object (or `C` for short). As the regression output shows, **salmon** automatically chooses a dummy/one-hot encoding of the levels and drops a baseline level. This behavior can be customized by specifying additional arguments to `C`. For example, **salmon** chose “1 Story” to be the baseline level for `Style`. To make “Other” the baseline level, we would specify the variable as `C("Style", baseline = "Other")`. Changing the baseline level does not affect the predictions from the model but does affect the interpretation of the coefficients.

Model 6 assumes no interaction between the factors, which is why the fitted lines in Figure 7 are parallel. To fit a model with an interaction term:

$$\begin{aligned}
 E[\log(\text{Price}(\$)) \mid \mathbf{x}] = & \beta_0 + \beta_1 I(\text{Style} = \text{"2 Story"}) + \\
 & \beta_2 I(\text{Style} = \text{"Other"}) + \beta_3 I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_4 I(\text{Style} = \text{"2 Story"}) I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_5 I(\text{Style} = \text{"Other"}) I(\text{Fire?} = \text{"Yes"})
 \end{aligned} \tag{7}$$

we use the `&` operator to include interactions of all orders between the two `Categorical` variables:

```
>>> house, fire = C("Style"), C("Fire?")
>>> interaction_model = LinearModel(house & fire, Log(Q("Price($)")))
>>> print(interaction_model)
```

```
log(Price($)) ~ 1 + Style + Fire? + (Style)(Fire?)
```

```
>>> interaction_model.fit(data)
```

	Coef.	SE	$t$	$p$	2.5%	97.5%
(Style{2 Story})(Fire?{Yes})	-0.0190	0.0289	-0.6593	0.5097	-0.0756	0.0376
(Style{Other})(Fire?{Yes})	-0.1638	0.0327	-5.008	5.813e-7	-0.2280	-0.0997
Style{2 Story}	0.1123	0.0217	5.179	2.385e-7	0.0698	0.1548
Style{Other}	-0.0740	0.0223	-3.318	9.181e-4	-0.1177	-0.0303
Fire?{Yes}	0.4347	0.0173	25.13	3e-126	0.4008	0.4687
Intercept	11.80	0.0119	987.7	0.000	11.78	11.82

Table 6: Coefficients and inferences for Model 7, `interaction_model`. (The labels have been abbreviated for space considerations.)

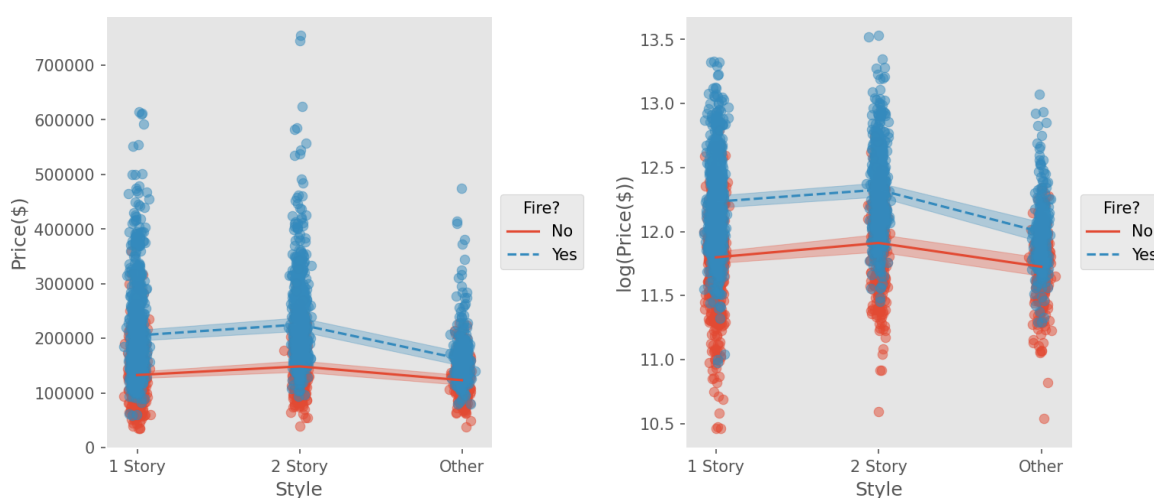


Figure 8: Two visualizations of Model 7, `interaction_model`, which contains two categorical predictors and their interaction. The plot on the left shows the original response (Sale Price), while the plot on the right shows the transformed response ( $\log(\text{Sale Price}(\$))$ ).

The same model would be specified in R as `log('Price($)' ) ~ 'Style' * 'Fire?'`. Note that R uses `*` to mean factor crossing, rather than multiplication.

```
>>> interaction_model.plot(confidence_band = alpha_val,
...   transformed_y_space = True)
```

In Figure 8, we see that the visualization of the fitted Model 7 is the classical interaction plot (Faraway 2016). Even when we allow for an interaction, the lines are still roughly parallel, suggesting that the interaction is weak.

Now, let's combine these categorical variables with the square footage into a single model.

	Coefficient	SE	<i>t</i>	<i>p</i>
(Style{2 Story})(Fire?{Yes})	-0.0389	0.0215	-1.807	0.0709
(Style{Other})(Fire?{Yes})	-0.1478	0.0244	-6.057	1.564e-9
Style{2 Story}	-0.1524	0.0171	-8.926	7.7e-19
Style{Other}	-0.1535	0.0167	-9.184	7.7e-20
Sq. Ft.	5.732e-4	1.188e-5	48.24	0.000
Fire?{Yes}	0.2160	0.0137	15.79	5.7e-54
Intercept	11.15	0.0161	691.6	0.000

Table 7: Coefficients and inferences for Model 8, `quant_cat_model`. (The confidence intervals are not shown for space considerations.)

For simplicity, we omit the quadratic term:

$$\begin{aligned}
 E[\log(\text{Price}(\$)) \mid \mathbf{x}] = & \beta_0 + \beta_1 I(\text{Style} = \text{"2 Story"}) + \\
 & \beta_2 I(\text{Style} = \text{"Other"}) + \beta_3 I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_4 I(\text{Style} = \text{"2 Story"}) I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_5 I(\text{Style} = \text{"Other"}) I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_6 (\text{Sq. Ft.})
 \end{aligned} \tag{8}$$

Adding the quantitative variable Sq. Ft. to the model is straightforward:

```

>>> house, fire, sqft = C("Style"), C("Fire?"), Q("Sq. Ft.")
>>> quant_cat_model = LinearModel((house & fire) + sqft, Log(Q("Price($)")))
>>> print(quant_cat_model)

```

```
log(Price($)) ~ 1 + Sq. Ft. + Style + Fire? + (Style)(Fire?)
```

```
>>> quant_cat_model.fit(data)
```

The same model would be specified in R by:

```
log('Price($)') ~ 'Sq. Ft.' + 'Style'*'Fire?'
```

```
>>> quant_cat_model.plot(confidence_band = alpha_val)
```

However, Model 8 assumed no interaction between the quantitative variable and the categorical variables. As one final step, we consider interacting the square footage with the categorical variables, which is equivalent to fitting a separate model for each possible combination of the




Figure 9: Visualization of the fitted Model 8, `quant_cat_model`. The quantitative variable is on the  $x$  axis, while each unique combination of the categorical variables is represented by a different line and color.

categorical variables:

$$\begin{aligned}
 E[\log(\text{Price}(\$)) \mid \mathbf{x}] = & \beta_0 + \beta_1(\text{Sq. Ft.}) + \beta_2 I(\text{Style} = \text{"2 Story"}) + \\
 & \beta_3 I(\text{Style} = \text{"Other"}) + \beta_4 I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_5 I(\text{Style} = \text{"2 Story"}) I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_6 I(\text{Style} = \text{"Other"}) I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_7 (\text{Sq. Ft.}) I(\text{Style} = \text{"2 Story"}) + \\
 & \beta_8 (\text{Sq. Ft.}) I(\text{Style} = \text{"Other"}) + \\
 & \beta_9 (\text{Sq. Ft.}) I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_{10} (\text{Sq. Ft.}) I(\text{Style} = \text{"2 Story"}) I(\text{Fire?} = \text{"Yes"}) + \\
 & \beta_{11} (\text{Sq. Ft.}) I(\text{Style} = \text{"Other"}) I(\text{Fire?} = \text{"Yes"})
 \end{aligned} \tag{9}$$

The sheer number of interaction terms would make this model difficult to specify, but because the variables are represented as composable objects in **salmon**, they can be stored in variables and reused, as illustrated below.

```

>>> house, fire, sqft = C("Style"), C("Fire?"), Q("Sq. Ft.")
>>> full_model = LinearModel(house & fire & sqft, Log(Q("Price($))))
>>> print(full_model)

log(Price($)) ~ 1 + Sq. Ft. + Style + Fire? + (Style)(Fire?) +
  (Sq. Ft.)(Style) + (Sq. Ft.)(Fire?) + (Sq. Ft.)(Style)(Fire?)

>>> full_model.fit(data)

```

	Coefficient	SE	<i>t</i>	<i>p</i>
(Sq. Ft.)(Style{2 Story})	-2.852e-4	4.556e-5	-6.260	4.4e-10
(Sq. Ft.)(Style{Other})	-4.985e-4	4.912e-5	-10.15	8.4e-24
(Style{2 Story})(Fire?{Yes})	-0.0766	0.0886	-0.8640	0.3877
(Style{Other})(Fire?{Yes})	-0.0236	0.0872	-0.2713	0.7862
Style{2 Story}	0.2222	0.0666	3.338	8.550e-4
Style{Other}	0.4554	0.0621	7.331	2.9e-13
(Sq. Ft.)(Style{2 Story})(Fire?Yes)	6.001e-5	5.605e-5	1.071	0.2844
(Sq. Ft.)(Style{Other})(Fire?Yes)	3.096e-5	6.166e-5	0.5021	0.6156
Fire?{Yes}	0.0592	0.0513	1.155	0.2482
Sq. Ft.	7.457e-4	2.922e-5	25.52	8e-130
(Sq. Ft.)(Fire?{Yes})	6.014e-5	3.821e-5	1.574	0.1156
Intercept	10.96	0.0341	321.1	0.000

Table 8: Coefficients and inferences for Model 9, `full_model`. (The confidence intervals are not shown for space considerations.)

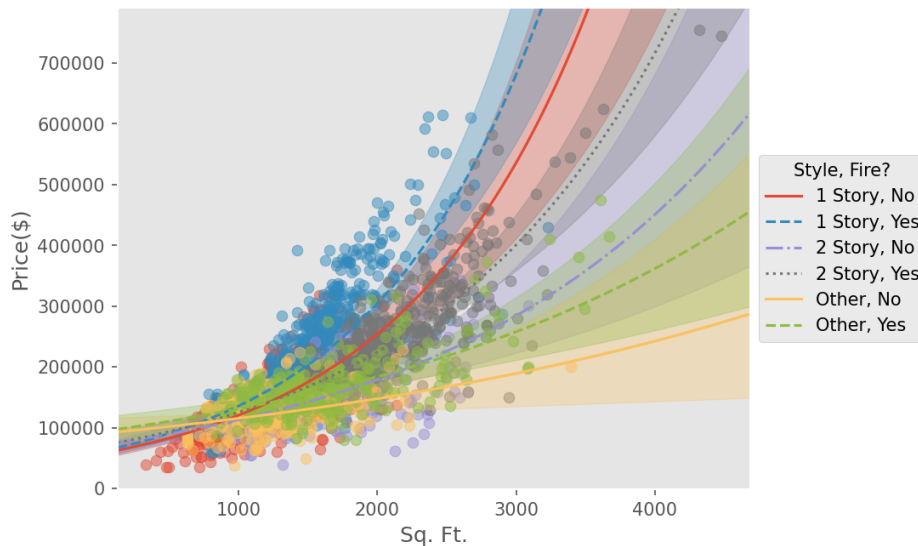


Figure 10: Visualization of the fitted Model 9, `full_model`, which includes an interaction between the quantitative and categorical variables.

```
>>> full_model.plot(confidence_band = alpha_val)
```

R supports a similarly concise syntax for the same model

```
R> log(`Price($)` ) ~ `Style` * `Fire?` * `Sq. Ft.`
```

but at the cost of making `*` mean something other than “multiplication”.

To use a regression model for prediction, we can call the `.predict()` method on new data. The new data must be a **pandas DataFrame** whose column names match the variable names in the model specification. In most cases, the new data will be in the same format as the

Predicted log(Price(\$))	
0	11.82
1	11.94
2	12.00
3	13.00
4	12.00

Table 9: First five predicted values from the `full_model` using the `DataFrame: new_data`.

	Predicted log(Price(\$))	2.5%	97.5%
0	11.82	11.36	12.29
1	11.94	11.48	12.41
2	12.00	11.54	12.47
3	13.00	12.53	13.47
4	12.00	11.53	12.46

Table 10: First five predicted values with 95% prediction intervals from the `full_model` using the `DataFrame: new_data`.

data used to fit the model, so this requirement will automatically be satisfied. In our example below, the new `DataFrame`, `new_data`, is simply the same as `data`.

```
>>> full_model.predict(new_data)
```

Accompanying confidence or prediction intervals can be obtained by passing the desired error rate  $\alpha$  to `confidence_interval` or `prediction_interval`:

```
>>> full_model.predict(new_data, prediction_interval = alpha_val)
```

Occasionally, it is useful to add variables to a regression model without estimating their coefficients. In **salmon**, this is achieved by simply shifting the response variable.

```
>>> offset = 3 * Q("Sq. Ft.")
>>> offset_model = LinearModel(C("Fire?"), Q("Price($)") - offset)
>>> print(offset_model)
```

```
Price($) - 3*(Sq. Ft.) ~ 1 + Fire?
```

Compare this with R, where an offset is specified using the `offset()` command:

```
R> SalePrice ~ Fire + offset(3 * SqFt)
```

Anecdotally, we have evidence that including an offset term by shifting the response is more intuitive. One of our colleagues, who uses R but was unfamiliar with the `offset()` function, attempted to implement offsets by shifting the response, which does not work in R. This motivated our decision to support shifting the response variable in **salmon**.

### 3. Model diagnostics

The previous section explained how to build regression models; this section focuses on how to evaluate them. There are both visual and analytical diagnostics; we only attempt to highlight a few examples of each. For a full description of the API, please refer to the online documentation.

In the following examples, we will check whether the assumptions are satisfied for Model 9, `full_model`, as well as compare it to Model 8, `quant_cat_model`, which is a nested model. Both models were defined in Section 2.

#### *Visual Diagnostics*

In Equation 1, we only made assumptions about the conditional expectation of the response  $y$ , given the explanatory variables  $\mathbf{x}$ . However, statistical inferences for linear regression require assumptions about the entire conditional distribution, not just the expectation. Perhaps the easiest way to specify this conditional distribution is to write

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_p x_p + \epsilon, \quad (10)$$

where the error term  $\epsilon$  is independent of  $\mathbf{x}$  and assumed to follow a  $\text{Normal}(0, \sigma^2)$  distribution. One can easily verify that this model is a special case of Model 1.

The assumption that the errors are independent and normally distributed with constant variance is usually assessed by inspecting the residuals. In **salmon**, the residual diagnostic plots can be generated directly from the fitted model:

```
>>> full_model.plot_residual_diagnostics()
```

Four plots of the residuals are produced, as shown in Figure 11: A normal Q-Q plot, a histogram, a scatterplot of the fitted values versus the residuals, and a line plot of the residuals versus their order in the data set. The residuals appear to be skewed to the left, but otherwise there do not appear to be obvious violations of the assumptions. The line plot of the residuals versus their order is not as useful for this data set, since the row ordering is arbitrary, but it could potentially reveal violations of independence in data sets where row order is significant, such as time series data.

To further investigate assumptions, we can produce residual plots and partial regression plots by calling the methods `model.residual_plots()` and `model.partial_plots()` respectively. These will each produce one plot per explanatory variable.

#### *Analytical Diagnostics*

Visual diagnostics are in the eye of the beholder, so analytical diagnostics are equally important. One common, if not always the most useful, diagnostic of a model's fit is the  $R^2$  value.

```
>>> full_model.r_squared()
```

```
0.652748378553996
```



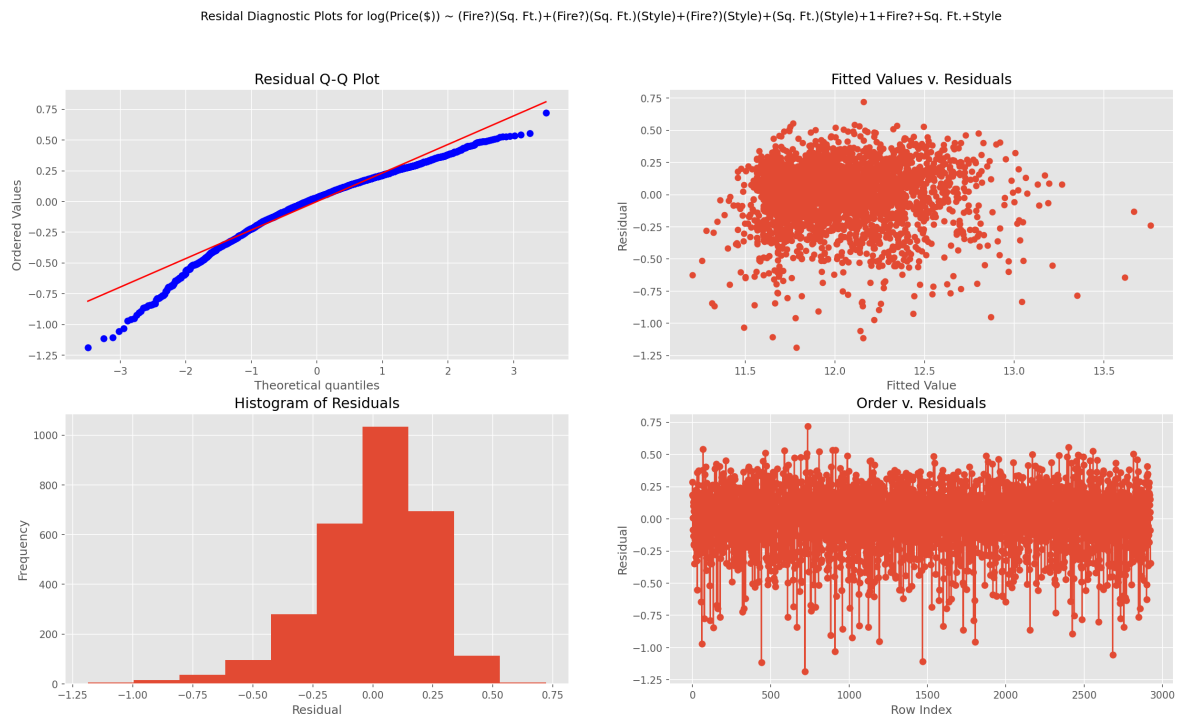


Figure 11: Residual diagnostic plots for the fitted regression model.

One problem with  $R^2$  is that it is monotonically increasing in the number of variables in the model. A better diagnostic is Adjusted- $R^2$ , which accounts for the number of variables. To calculate Adjusted- $R^2$ , we specify the argument `adjusted = True`. Besides  $R^2$ , AIC and BIC are also available.

A more formal way of evaluating a model is to test it against another model. The `anova` command, when called on a single model, returns the results of the omnibus  $F$  test, as well as the partial  $F$  test for the each individual variable in the model.

```
>>> from salmon import anova
>>> anova(full_model)
```

The inclusion of interactions between the quantitative variable (square footage) and the categorical variables (fireplace and house style) adds complexity to the model that may not be supported by the data. To test this, we can perform a partial  $F$  test comparing the full model to a reduced model without these interaction terms using the `anova` command again, except passing in two fitted models. This mirrors the behavior of the `anova()` function in R.

```
>>> anova(full_model, quant_cat_model)
```

The results of the partial  $F$  test confirm that interactions significantly improve the fit of the model to the data.

	DF	SS Err.	SS Reg.	$F$	$p$
Global Test	11	309.578	309.578	500.765	0.000
(Style)(Sq. Ft.)	2	169.819	303.415	54.828	4.2e-24
(Fire?)(Style)	2	163.699	309.536	0.373	0.6883
Style	2	166.776	306.458	27.755	1.1e-12
(Fire?)(Style)(Sq. Ft.)	2	163.721	309.513	0.574	0.5630
Fire?	1	163.731	309.503	1.333	0.2482
Sq. Ft.	1	200.263	272.971	651.347	8e-130
(Fire?)(Sq. Ft.)	1	163.796	309.438	2.477	0.1156
Error	2912				

Table 11: Global  $F$  test results for `full_model`.

	DF	SS Err.	SS Reg.	$F$	$p$
Full Model	11	163.657	309.578		
Reduced Model	5	179.015	294.219	54.655	2.0e-54
Error	2912				

Table 12: Partial  $F$  test results from comparing `full_model` to `quant_cat_model`.

## 4. Automatic model building

The **salmon** package also provides functions to automate the model building process. In the following examples, we use **salmon** to automatically select the “best” model, according to a metric of our choosing, from among the following variables:  $\text{Sq. Ft.} - E[\text{Sq. Ft.}]$ ,  $(\text{Sq. Ft.} - E[\text{Sq. Ft.}])^2$ ,  $I(\text{Style} = \text{“2 Story”})$ ,  $I(\text{Style} = \text{“Other”})$ ,  $I(\text{Fire?} = \text{“Yes”})$ , plus interactions between all categorical and quantitative variables. The response variable in this scenario is  $\log(\text{Price}(\$))$ .

To utilize the model building features of **salmon**, we first specify a model that contains all of the variables under consideration:

```
>>> quant_vars = Poly(Cen(Q("Sq. Ft.")), 2)
>>> all_terms_model = LinearModel(
...     (1 + C("Style")) * (1 + C("Fire?")) * (1 + quant_vars),
...     Log(Q("Price($))))
>>> print(all_terms_model)
```

```
log(Sale Price($)) ~ 1 + (Sq. Ft. - E(Sq. Ft.)) +
(Sq. Ft. - E(Sq. Ft.))^2 + Fire? +
(Fire?)(Sq. Ft. - E(Sq. Ft.)) +
(Fire?)((Sq. Ft. - E(Sq. Ft.))^2) +
Style + (Style)(Sq. Ft. - E(Sq. Ft.)) +
(Style)((Sq. Ft. - E(Sq. Ft.))^2) +
(Fire?)(Style) +
(Fire?)(Style)(Sq. Ft. - E(Sq. Ft.)) +
(Fire?)(Style)((Sq. Ft. - E(Sq. Ft.))^2)
```

Stepwise selection methods are supported through the `stepwise` function. Forward stepwise is specified by `forward = True`, while backward stepwise is specified by `forward = False`. Additionally, we specify what metric to optimize (e.g., AIC, BIC,  $R^2$ ) using the `method` parameter. For example, if we use forward stepwise to optimize for BIC, we obtain the following “best” model:

```
>>> from salmon import stepwise
>>> results = stepwise(full_model = all_terms_model, metric_name = "BIC",
...     forward = True, data = data)
>>> print(results["metric"])
```

```
BIC | -168.69617383126013
```

```
>>> print(results["best_model"])
```

```
log(Price($)) ~ 1 +
  (Fire?)(Sq. Ft. - E(Sq. Ft.)) +
  (Style)(Sq. Ft. - E(Sq. Ft.)) +
  Fire? + Style + (Sq. Ft. - E(Sq. Ft.)) +
  (Sq. Ft. - E(Sq. Ft.))^2 +
  (Fire?)(Style)(Sq. Ft. - E(Sq. Ft.))^2
```

The procedure produced a model associated with the lowest BIC score, but at the cost of violating what some would consider best modeling practices: namely, the inclusion of interactions and higher-order terms without main effects and lower-order terms (Cox 1984). For example,  $(\text{Fire?})(\text{Style})(\text{Sq. Ft.} - E(\text{Sq. Ft.}))^2$  is included, but not  $(\text{Fire?})(\text{Style})$  or  $(\text{Fire?})(\text{Style})(\text{Sq. Ft.} - E(\text{Sq. Ft.}))$ .

Due the object-oriented nature of **salmon**, the package can keep track of inter-variable relationships. To employ this feature, we simply specify the parameter `naive = False` when calling the `stepwise` function:

```
>>> results = stepwise(full_model = all_terms_model, metric_name = "BIC",
...     forward = True, naive = False, data = data)
>>> print(results["metric"])
```

```
BIC | -187.50648709382568
```

```
>>> print(results["best_model"])
```

```
log(Price($)) ~ 1 + (Fire?)(Sq. Ft. - E(Sq. Ft.)) +
  (Fire?)(Sq. Ft. - E(Sq. Ft.))^2 +
  (Style)(Sq. Ft. - E(Sq. Ft.)) +
  (Style)(Sq. Ft. - E(Sq. Ft.))^2 +
  Fireplace? + Style + (Sq. Ft. - E(Sq. Ft.)) +
  (Sq. Ft. - E(Sq. Ft.))^2
```

Model	Task	<b>salmon</b>	<b>statsmodels</b>	R (1m)
Simple Linear Regression (1 quantitative predictor)	fit	2.759	8.701	1.505
	predict	0.607	2.267	0.350
Multiple Regression (4 quant. + 1 cat. predictors)	fit	3.637	15.590	5.098
	predict	1.421	6.536	1.020
Multiple Regression (2 quant. + 1 cat. + interactions)	fit	7.810	11.576	2.344
	predict	2.671	3.994	1.171
Multiple Regression (1000 quant.)	fit	2198.1	7491.4	5280.5
	predict	215.73	22.80	191.41

Table 13: Timing comparisons of **salmon** and competitors on four models. For each model, we time two tasks: 1) fitting the model (including inferential statistics) and 2) using the fitted model to make predictions. All reported times are in milliseconds.

Now, the selected model is guaranteed to not violate standard model building practices. For example, if an interaction term is included, the main effects will be also.

## 5. Integration with Python ecosystem

The **salmon** package is deliberately built on top of **pandas**, in order to leverage the power of its **DataFrame** for organizing heterogeneous data (McKinney 2010). It expects **pandas DataFrames** as inputs, and it outputs results as **pandas DataFrames**. However, **salmon** is also integrated with other packages within the Python data science ecosystem.

First, the ‘**LinearModel**’ class in **salmon** implements `.fit()` and `.predict()` methods, so any ‘**LinearModel**’ can be used as **scikit-learn** estimators. This means that all of **scikit-learn**’s routines for model selection and evaluation, such as cross validation, will work with models that are created in **salmon** (Buitinck *et al.* 2013).

Second, **salmon** plots are produced using **matplotlib**. These plots can be further edited and customized in **matplotlib** because every plotting command in **salmon** returns a **matplotlib Figure** (Hunter 2007).

## 6. Timing and numerical comparisons

We conducted timing comparisons between **salmon**, **statsmodels**, and R. The results are shown in Table 13. The **salmon** package is consistently faster than **statsmodels** in all scenarios. It is also faster than R on large examples (with  $p = 1000$  predictors), although it is still slower than R on small examples, perhaps because formulas are native in R. The speed of **salmon** on large examples is likely due to the careful handling of the linear algebra, some details of which are described in Appendix A.2.

We also compare the numerical stability of **salmon**, **statsmodels**, and R on a synthetic example. First, we generate a grid of  $n = 1000$  evenly-spaced values in the range  $[10^8 - 1, 10^8 + 1]$  for the

	<b>salmon</b>	<b>statsmodels</b>	R (1m)
Coefficient	2.095 (0.1059)	2.000 ( $3 \times 10^{-10}$ )	NA (NA)
Intercept	3046697 (5372320)	0.00000002 ( $3 \times 10^{-18}$ )	200000002 (0.0475)

Table 14: Estimated parameters (with standard errors in parentheses) from **salmon** and competitors on synthetic data simulated according to Equation 11. Only **salmon** recovers the correct coefficients. R fails to produce a coefficient at all, while **statsmodels** returns the incorrect coefficients with a high degree of confidence.

explanatory variable  $x$ . Then, we generate the response variable  $y$  according to the relation

$$y_i = 2 + 2x_i + \epsilon_i, \quad (11)$$

where  $\epsilon_1, \dots, \epsilon_n$  are i.i.d.  $\text{Normal}(0, 1)$ . Model 11 describes a simple linear regression model with  $\beta_0 = \beta_1 = 2$ . The large values of  $x$  make it nearly collinear with the intercept.

The estimated parameters ( $\hat{\beta}_0$  and  $\hat{\beta}_1$ ) from the three packages are shown in Table 14. R fails to estimate a coefficient at all. The behavior of **statsmodels** is more troubling; it throws a warning, but proceeds to report incorrect values with a high degree of precision. Only **salmon** returns the correct values, due to careful handling of numerical linear algebra (see Appendix A.2 for more details). Note that although the intercept is suspiciously large, this is because the standard error of the intercept is very large. Recall that the standard deviation of the intercept is  $\sigma \sqrt{\frac{1}{n} + \frac{\bar{x}^2}{\sum_i (x_i - \bar{x})^2}} \approx 5471751$  for the parameters of our simulation.

The **Jupyter** notebooks to reproduce these experiments can be found in the supplementary materials and on GitHub at [https://github.com/ajboyd2/salmon/tree/master/paper\\_outputs](https://github.com/ajboyd2/salmon/tree/master/paper_outputs).

## 7. Conclusion

The **salmon** package allows for convenient and intuitive model building, using, and diagnosing for linear regression within Python. As our examples have demonstrated, the consistent interface allows for an linear regression experience in Python that reproduces – and in some cases, enhances – the experience in R.

Future work in **salmon** will pertain to extending the package’s capabilities by implementing other types of models beyond linear regression, such as generalized linear models, as well as providing additional tests that are supported in R, such as linear contrasts.

## Acknowledgments

We are indebted to Ben Mangelsdorf, who contributed code to the **salmon** package and helped with the preparation of this paper. We are also grateful to Fernando Perez, Nicholas Russo, Jonathan Taylor, and John Walker for their feedback, as well as to the Bill and Linda Frost Fund for support.

## References

- Buitinck L, Louppe G, Blondel M, Pedregosa F, Mueller A, Grisel O, Niculae V, Prettenhofer P, Gramfort A, Grobler J, Layton R, Vanderplas J, Joly A, Holt B, Varoquaux G (2013). “API Design for Machine Learning Software: Experiences from the **scikit-learn** Project.” In *European Conference on Machine Learning and Principles and Practices of Knowledge Discovery in Databases Workshop: Languages for Data Mining and Machine Learning*, pp. 108–122.
- Chatterjee S, Hadi AS (2015). *Regression Analysis by Example*. John Wiley & Sons.
- Cox DR (1984). “Interaction.” *International Statistical Review*, pp. 1–24. doi:10.2307/1403235.
- De Cock D (2011). “Ames, Iowa: Alternative to the Boston Housing Data as an End of Semester Regression Project.” *Journal of Statistics Education*, **19**(3). doi:10.1080/10691898.2011.11889627.
- Faraway JJ (2016). *Linear Models with R*. Chapman & Hall/CRC.
- Hunter JD (2007). “**matplotlib**: A 2D Graphics Environment.” *Computing in Science & Engineering*, **9**(3), 90–95. doi:10.1109/mcse.2007.55.
- McKinney W (2010). “Data Structures for Statistical Computing in Python.” In S Van der Walt, J Millman (eds.), *Proceedings of the 9th Python in Science Conference*, pp. 51–56.
- Minitab Inc (2017). *Minitab 17: Statistical Software for Process Control, Improvement and Education*. Minitab Inc., State College. URL <https://www.minitab.com/>.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, et al. (2011). “**scikit-learn**: Machine Learning in Python.” *Journal of Machine Learning Research*, **12**(85), 2825–2830.
- R Core Team (2023). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Ross K, Sun DL (2019). “**symbulate**: Simulation in the Language of Probability.” *Journal of Statistics Education*, pp. 1–17. doi:10.1080/10691898.2019.1600387.
- SAS Institute Inc (2019). *JMP, Version 15*. Cary. URL <https://www.JMP.com/>.
- Seabold S, Perktold J (2010). “**statsmodels**: Econometric and Statistical Modeling with Python.” In *Proceedings of the 9th Python in Science Conference*, volume 57, p. 61.
- Smith NJ, et al. (2022). **patsy**. Python package version 0.5.3, URL <https://pypi.org/project/patsy/>.
- The MathWorks Inc (2021). *MATLAB – The Language of Technical Computing, Version R2021a*. Natick. URL <https://www.mathworks.com/products/matlab/>.
- Van Rossum G, et al. (2011). *Python Programming Language*. URL <https://www.python.org/>.

## A. Implementation details

### A.1. Expression tree

A goal for **salmon** was to be able to represent variables in an object oriented fashion that also allowed easy manipulations/transformations. The solution to this was to represent the variables as an abstract syntax tree, borrowing heavily from a programming languages approach. Every expression of variables can be modeled as a tree where each inner node is an operation such as addition, multiplication, or transformation and each leaf is a variable (either quantitative or categorical). The following code and figure are equivalent in representation.

```
>>> Q("a") - C("b") * C("c") + Q("d") ** 3 * Q("e") +
...   Sin(Q("f") ** 5) * Log(Q("g") + Q("h") + 1)
```

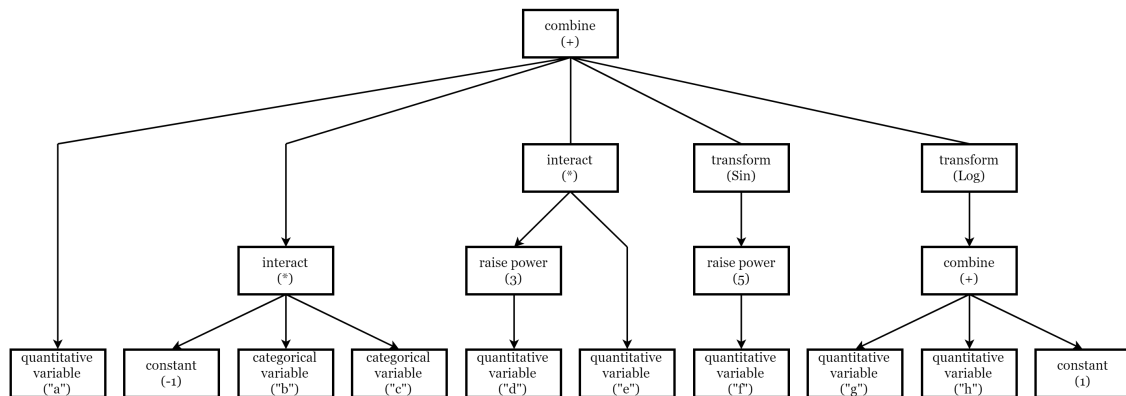


Figure 12: Expression tree.

This design allowed for a recursive style implementation when defining actions amongst variables. We denote the class ‘**Expression**’ to denote the most general and all encompassing structure of variables. The exact object composition scheme can be described in BNF with the following grammar:

```

<Expression> ::= <Variable>
               | <Constant>
               | <Expression> + <Expression>
               | <Expression> - <Expression>
               | <Expression> * <Expression>
               | <Expression> / <Expression>
               | <Expression> ** <Constant>

<Variable>   ::= Var(<Name>)
               | <Quantitative>
               | <Categorical>
               | <Transformation>(<Expression>)
  
```

```

| <Variable> ** <Constant>
| <Variable> * <Constant>
| <Variable> / <Constant>
| <Variable> * <Variable>
| <Variable> / <Variable>

<Quantitative> ::= Q(<Name>)

<Categorical> ::= C(<Name>)

<Transformation> ::= Sin
| Cos
| Exp
| Log
| Log10
| Standardize
| Center
| Identity

<Constant> ::= int
| float
| <Constant> + <Constant>
| <Constant> - <Constant>
| <Constant> * <Constant>
| <Constant> / <Constant>

<Name> ::= str

```

## A.2. Model fitting

Our general strategy for fitting the linear regression model is similar to R's (R Core Team 2023). We first calculate a QR decomposition of the design matrix  $X$ , which allows us to calculate the regression coefficients  $\hat{\beta}$  by solving the system  $R\hat{\beta} = Q^\top \mathbf{y}$  and the estimated covariance matrix  $\hat{\Sigma} = \hat{\sigma}^2(X^\top X)^{-1}$  by solving the system  $R^\top R\hat{\Sigma} = I$  (since  $X^\top X = R^\top R$ ).

When the model includes an intercept, R calculates the coefficients by appending a column of ones to the design matrix

$$\tilde{X} = \begin{pmatrix} \mathbf{1} & X \end{pmatrix}.$$

The regression coefficients are then given by

$$\hat{\beta} := \begin{pmatrix} \hat{\beta}_0 \\ \hat{\beta}_{1:p} \end{pmatrix} = (\tilde{X}^\top \tilde{X})^{-1} \tilde{X}^\top \mathbf{y}.$$

We take a slightly different approach, centering all of the variables first. The centered versions of the design matrix  $X$  and the response variable  $\mathbf{y}$  are given by

$$\begin{aligned} X_c &= X - \frac{1}{n} \mathbf{1}\mathbf{1}^\top X & \mathbf{y}_c &= \mathbf{y} - \frac{1}{n} \mathbf{1}\mathbf{1}^\top \mathbf{y} \\ &:= X - \mathbf{1}\bar{\mathbf{x}}^\top & &:= \mathbf{y} - \mathbf{1}\bar{y} \end{aligned}$$



It is well-known that centering does not change the coefficients of the explanatory variables (Chatterjee and Hadi 2015). That is, we can calculate  $\hat{\beta}_{1:p}$  as

$$\hat{\beta}_{1:p} = (X_c^\top X_c)^{-1} X_c^\top \mathbf{y}_c$$

and recover the intercept as

$$\hat{\beta}_0 = \mathbf{y} - \bar{\mathbf{x}}^\top \hat{\beta}_{1:p}.$$

The advantage of centering is that the computations are more numerically stable.

To calculate the covariance matrix,  $\hat{\Sigma}$ , the main hurdle is computing  $(\tilde{X}^\top \tilde{X})^{-1}$ . We compute this without explicitly constructing  $\tilde{X}$ , using properties of block matrices. To our knowledge, these formulas do not appear in the literature, so we reproduce the derivations here for completeness.

First, observe that the inverse of  $\tilde{X}^\top \tilde{X}$  can be written in block form as

$$(\tilde{X}^\top \tilde{X})^{-1} = \begin{pmatrix} n & \mathbf{1}^\top X \\ X^\top \mathbf{1} & X^\top X \end{pmatrix}^{-1} = \begin{pmatrix} a & \mathbf{b}^\top \\ \mathbf{b} & C \end{pmatrix}.$$

Next,  $C$  can be obtained as the inverse of the Schur complement of the block matrix  $\tilde{X}^\top \tilde{X}$ :

$$\begin{aligned} C &= \left( X^\top X - (X^\top \mathbf{1}) \frac{1}{n} (\mathbf{1}^\top X) \right)^{-1} \\ &= \left( X^\top (I - P) X \right)^{-1} && (P := \frac{1}{n} \mathbf{1} \mathbf{1}^\top \text{ is a projection matrix}) \\ &= \left( X^\top (I - P)^2 X \right)^{-1} && ((I - P)^2 = I - P) \\ &= \left( ((I - P) X)^\top (I - P) X \right)^{-1} && (I - P \text{ is symmetric}) \\ &= (X_c^\top X_c)^{-1}. \end{aligned}$$

Finally, we compute  $a$  and  $\mathbf{b}$  in terms of  $C$ .

$$\begin{aligned} \mathbf{b} &= -C (X^\top \mathbf{1}) \frac{1}{n} \\ &= -C \bar{\mathbf{x}} \\ a &= \frac{1}{n} + \frac{1}{n} (\mathbf{1}^\top X) C (X^\top \mathbf{1}) \frac{1}{n} \\ &= \frac{1}{n} - \bar{\mathbf{x}}^\top \mathbf{b}. \end{aligned}$$

**Affiliation:**

Dennis L. Sun  
California Polytechnic State University  
Department of Statistics  
1 Grand Ave  
San Luis Obispo, CA 93407, United States of America  
E-mail: [dsun09@calpoly.edu](mailto:dsun09@calpoly.edu)  
URL: <https://dlsun.github.io/>