Example 36. (warmup) $\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = x_1 \begin{bmatrix} 1 \\ 3 \\ 0 \end{bmatrix} + x_2 \begin{bmatrix} 2 \\ 1 \\ 5 \end{bmatrix}$

Note that this means that the system of equations $\begin{array}{rcrcl} x_1 & + & 2x_2 & = & 1 \\ 3x_1 & + & x_2 & = & 1 \\ & & 5x_2 & = & 1 \end{array}$ can also be written as $\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$.

[This was the motivation for introducing matrix-vector multiplication.]

In the same way, any system can be written as $A\boldsymbol{x} = \boldsymbol{b}$, where A is a matrix and $\boldsymbol{b}$ a vector.
In particular, this makes it obvious that:

$$A\boldsymbol{x} = \boldsymbol{b} \text{ is consistent} \iff \boldsymbol{b} \text{ is in } \operatorname{col}(A)$$

Recall that, by the FTLA, $\operatorname{col}(A)$ and $\operatorname{null}(A^T)$ are orthogonal complements.

Theorem 37. $A\boldsymbol{x} = \boldsymbol{b}$ is consistent $\iff$ $\boldsymbol{b}$ is orthogonal to $\operatorname{null}(A^T)$

Proof. $A\boldsymbol{x} = \boldsymbol{b}$ is consistent $\iff$ $\boldsymbol{b}$ is in $\operatorname{col}(A)$ $\overset{\text{FTLA}}{\iff}$ $\boldsymbol{b}$ is orthogonal to $\operatorname{null}(A^T)$

Note. $\boldsymbol{b}$ is orthogonal to $\operatorname{null}(A^T)$ means that $\boldsymbol{y}^T\boldsymbol{b} = 0$ whenever $\boldsymbol{y}^T A = \boldsymbol{0}$. Why?!

Example 38. Let $A = \begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix}$. For which $\boldsymbol{b}$ does $A\boldsymbol{x} = \boldsymbol{b}$ have a solution?

Solution. (old)

$$\left[\begin{array}{cc|c} 1 & 2 & b_1 \\ 3 & 1 & b_2 \\ 0 & 5 & b_3 \end{array}\right] \overset{R_2 - 3R_1 \Rightarrow R_2}{\rightsquigarrow} \left[\begin{array}{cc|c} 1 & 2 & b_1 \\ 0 & -5 & -3b_1 + b_2 \\ 0 & 5 & b_3 \end{array}\right] \overset{R_3 + R_2 \Rightarrow R_3}{\rightsquigarrow} \left[\begin{array}{cc|c} 1 & 2 & b_1 \\ 0 & -5 & -3b_1 + b_2 \\ 0 & 0 & -3b_1 + b_2 + b_3 \end{array}\right]$$

So, $A\boldsymbol{x} = \boldsymbol{b}$ is consistent if and only if $-3b_1 + b_2 + b_3 = 0$.

Solution. (new) We determine a basis for $\operatorname{null}(A^T)$:

$$\begin{bmatrix} 1 & 3 & 0 \\ 2 & 1 & 5 \end{bmatrix} \overset{R_2 - 2R_1 \Rightarrow R_2}{\rightsquigarrow} \begin{bmatrix} 1 & 3 & 0 \\ 0 & -5 & 5 \end{bmatrix} \overset{-\frac{1}{5}R_2 \Rightarrow R_2}{\rightsquigarrow} \begin{bmatrix} 1 & 3 & 0 \\ 0 & 1 & -1 \end{bmatrix} \overset{R_1 - 3R_2 \Rightarrow R_1}{\rightsquigarrow} \begin{bmatrix} 1 & 0 & 3 \\ 0 & 1 & -1 \end{bmatrix}$$

We read off from the RREF that $\operatorname{null}(A^T)$ has basis $\begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix}$.

$\boldsymbol{b}$ has to be orthogonal to $\operatorname{null}(A^T)$. That is, $\boldsymbol{b} \cdot \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix} = 0$. As above!

Comment. Below is how we can use Sage to (try and) solve $A\boldsymbol{x} = \boldsymbol{b}$ for $\boldsymbol{b} = \begin{bmatrix} 1 \\ 1 \\ 2 \end{bmatrix}$ and $\boldsymbol{b} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$.

```
Sage] A = matrix([[1,2],[3,1],[0,5]])

Sage] A.solve_right(vector([1,1,2]))
```

$$\left(\frac{1}{5}, \frac{2}{5}\right)$$

```
Sage] A.solve_right(vector([1,1,1]))

   Traceback (most recent call last):
   ValueError: matrix equation has no solutions
```

Armin Straub
straub@southalabama.edu

Least squares

Example 39. Not all linear systems have solutions.

In fact, for many applications, data needs to be fitted and there is no hope for a perfect match.

For instance, $A\boldsymbol{x} = \boldsymbol{b}$ with

$$\begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix} \boldsymbol{x} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

has no solution:

- $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$ is not in $\text{col}(A)$ since $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \cdot \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix} \neq 0$ (see previous example).
- Instead of giving up, we want the $\boldsymbol{x}$ which makes $A\boldsymbol{x}$ and $\boldsymbol{b}$ as close as possible.
- Such $\boldsymbol{x}$ is characterized by the error $A\boldsymbol{x} - \boldsymbol{b}$ being **orthogonal** to $\text{col}(A)$ (i.e. all possible $A\boldsymbol{x}$).

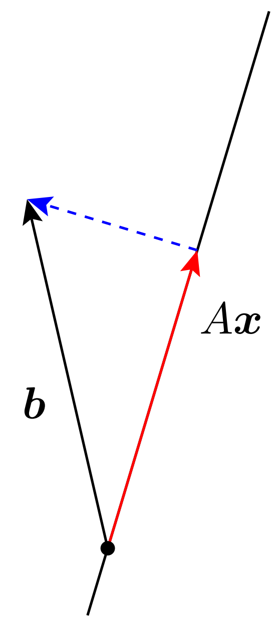

Definition 40. $\hat{\boldsymbol{x}}$ is a **least squares solution** of the system $A\boldsymbol{x} = \boldsymbol{b}$ if $\hat{\boldsymbol{x}}$ is such that $A\hat{\boldsymbol{x}} - \boldsymbol{b}$ is as small as possible (i.e. minimal norm).

- If $A\boldsymbol{x} = \boldsymbol{b}$ is consistent, then $\hat{\boldsymbol{x}}$ is just an ordinary solution. (in that case, $A\hat{\boldsymbol{x}} - \boldsymbol{b} = \boldsymbol{0}$)
- Interesting case: $A\boldsymbol{x} = \boldsymbol{b}$ is inconsistent. (in particular, if the system is overdetermined)

The normal equations

The following result provides a straightforward recipe (thanks to the FTLA) to find least squares solutions for all systems $A\boldsymbol{x} = \boldsymbol{b}$.

Theorem 41. $\hat{\boldsymbol{x}}$ is a least squares solution of $A\boldsymbol{x} = \boldsymbol{b}$
$\iff A^T A\hat{\boldsymbol{x}} = A^T\boldsymbol{b}$ (the **normal equations**)

Proof.

$\hat{\boldsymbol{x}}$ is a least squares solution of $A\boldsymbol{x} = \boldsymbol{b}$

$\iff A\hat{\boldsymbol{x}} - \boldsymbol{b}$ is as small as possible

$\iff A\hat{\boldsymbol{x}} - \boldsymbol{b}$ is orthogonal to $\text{col}(A)$

$\overset{\text{FTLA}}{\iff} A\hat{\boldsymbol{x}} - \boldsymbol{b}$ is in $\text{null}(A^T)$

$\iff A^T(A\hat{\boldsymbol{x}} - \boldsymbol{b}) = \boldsymbol{0}$

$\iff A^T A\hat{\boldsymbol{x}} = A^T\boldsymbol{b}$ □

Example 42. Find the least squares solution to $A\boldsymbol{x} = \boldsymbol{b}$, where

$$A = \begin{bmatrix} 1 & 1 \\ -1 & 1 \\ 0 & 0 \end{bmatrix}, \quad \boldsymbol{b} = \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix}.$$

Solution. First, $A^T A = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 & 1 \\ -1 & 1 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$ and $A^T \boldsymbol{b} = \begin{bmatrix} 1 & -1 & 0 \\ 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$.

Hence, the normal equations $A^T A \hat{\boldsymbol{x}} = A^T \boldsymbol{b}$ take the form $\begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix} \hat{\boldsymbol{x}} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$.
Solving, we immediately find $\hat{\boldsymbol{x}} = \begin{bmatrix} 1/2 \\ 3/2 \end{bmatrix}$.

Check. Since $A\hat{\boldsymbol{x}} = \begin{bmatrix} 2 \\ 1 \\ 0 \end{bmatrix}$, the error is $A\hat{\boldsymbol{x}} - \boldsymbol{b} = \begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix}$. Recall that the error must be orthogonal to $\operatorname{col}(A)$!

This error is indeed orthogonal to $\operatorname{col}(A)$ because $\begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ -1 \\ 0 \end{bmatrix} = 0$ and $\begin{bmatrix} 0 \\ 0 \\ -1 \end{bmatrix} \cdot \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix} = 0$.

Comment. Why are the normal equations so particularly simple (compare with example below for the typical case) here? Note how each entry of the product $A^T A$ is computed as the dot product of two columns of A (matrix products of a row of A^T times a column of A). That $A^T A$ is a diagonal matrix reflects the fact that the two columns of A are orthogonal to each other.

Example 43. Find the least squares solution to $A\boldsymbol{x} = \boldsymbol{b}$, where

$$A = \begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix}, \quad \boldsymbol{b} = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}.$$

Solution. First, $A^T A = \begin{bmatrix} 1 & 3 & 0 \\ 2 & 1 & 5 \end{bmatrix} \begin{bmatrix} 1 & 2 \\ 3 & 1 \\ 0 & 5 \end{bmatrix} = \begin{bmatrix} 10 & 5 \\ 5 & 30 \end{bmatrix}$ and $A^T \boldsymbol{b} = \begin{bmatrix} 1 & 3 & 0 \\ 2 & 1 & 5 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 4 \\ 8 \end{bmatrix}$.

Hence, the normal equations $A^T A \hat{\boldsymbol{x}} = A^T \boldsymbol{b}$ take the form $\begin{bmatrix} 10 & 5 \\ 5 & 30 \end{bmatrix} \hat{\boldsymbol{x}} = \begin{bmatrix} 4 \\ 8 \end{bmatrix}$.

Since $\begin{bmatrix} 10 & 5 \\ 5 & 30 \end{bmatrix}^{-1} = \frac{1}{275} \begin{bmatrix} 30 & -5 \\ -5 & 10 \end{bmatrix} = \frac{1}{55} \begin{bmatrix} 6 & -1 \\ -1 & 2 \end{bmatrix}$, we find $\hat{\boldsymbol{x}} = \frac{1}{55} \begin{bmatrix} 6 & -1 \\ -1 & 2 \end{bmatrix} \begin{bmatrix} 4 \\ 8 \end{bmatrix} = \frac{1}{55} \begin{bmatrix} 16 \\ 12 \end{bmatrix}$.

Check. Since $A\hat{\boldsymbol{x}} = \frac{1}{55} \begin{bmatrix} 40 \\ 60 \\ 60 \end{bmatrix}$, the error $A\hat{\boldsymbol{x}} - \boldsymbol{b} = \frac{1}{55} \begin{bmatrix} -15 \\ 5 \\ 5 \end{bmatrix} = \frac{1}{11} \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix}$ must be orthogonal to $\operatorname{col}(A)$.

The error is indeed orthogonal to $\operatorname{col}(A)$ because $\begin{bmatrix} 1 \\ 3 \\ 0 \end{bmatrix} \cdot \frac{1}{11} \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix} = 0$ and $\begin{bmatrix} 2 \\ 1 \\ 5 \end{bmatrix} \cdot \frac{1}{11} \begin{bmatrix} -3 \\ 1 \\ 1 \end{bmatrix} = 0$.

Example 44. Find the least squares solution to $A\boldsymbol{x}=\boldsymbol{b}$, where

$$A=\begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}, \quad \boldsymbol{b}=\begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}.$$

Solution. First, $A^TA=\begin{bmatrix} 4 & 0 & 1 \\ 0 & 2 & 1 \end{bmatrix}\begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}=\begin{bmatrix} 17 & 1 \\ 1 & 5 \end{bmatrix}$ and $A^T\boldsymbol{b}=\begin{bmatrix} 4 & 0 & 1 \\ 0 & 2 & 1 \end{bmatrix}\begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}=\begin{bmatrix} 19 \\ 11 \end{bmatrix}$.

Hence, the normal equations $A^TA\hat{\boldsymbol{x}}=A^T\boldsymbol{b}$ take the form $\begin{bmatrix} 17 & 1 \\ 1 & 5 \end{bmatrix}\hat{\boldsymbol{x}}=\begin{bmatrix} 19 \\ 11 \end{bmatrix}$.

Solving, we immediately find $\hat{\boldsymbol{x}}=\begin{bmatrix} 1 \\ 2 \end{bmatrix}$.

Check. The error $A\hat{\boldsymbol{x}}-\boldsymbol{b}=\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix}$ is indeed orthogonal to $\operatorname{col}(A)$. Because $\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix}\cdot\begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}=0$ and $\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix}\cdot\begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix}=0$.

Application: least squares lines

Given data points (x_i,y_i), we wish to find optimal parameters a,b such that $y_i\approx a+bx_i$ for all i.

Example 45. Determine the line that "best fits" the data points $(2,1),(5,2),(7,3),(8,3)$.

Comment. Can you see that there is no line fitting the data perfectly? (Check out the last two points!)

Solution. We need to determine the values a,b for the best-fitting line $y=a+bx$.
If there was a line that fit the data perfectly, then:

$$\begin{aligned} a+2b&=1 && (2,1)\\ a+5b&=2 && (5,2)\\ a+7b&=3 && (7,3)\\ a+8b&=3 && (8,3)\end{aligned}$$

In matrix form, this is: $\underbrace{\begin{bmatrix} 1 & x_1 \\ 1 & x_2 \\ 1 & x_3 \\ 1 & x_4 \end{bmatrix}}_{\text{design matrix } X}\begin{bmatrix} a \\ b \end{bmatrix}=\underbrace{\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix}}_{\text{observation vector } \boldsymbol{y}}$ (writing the points as (x_i,y_i))

Using our points, these equations become $\begin{bmatrix} 1 & 2 \\ 1 & 5 \\ 1 & 7 \\ 1 & 8 \end{bmatrix}\begin{bmatrix} a \\ b \end{bmatrix}=\begin{bmatrix} 1 \\ 2 \\ 3 \\ 3 \end{bmatrix}$. [This system is inconsistent (as expected).]

We compute a least squares solution.

$$X^TX=\begin{bmatrix} 1 & 1 & 1 & 1 \\ 2 & 5 & 7 & 8 \end{bmatrix}\begin{bmatrix} 1 & 2 \\ 1 & 5 \\ 1 & 7 \\ 1 & 8 \end{bmatrix}=\begin{bmatrix} 4 & 22 \\ 22 & 142 \end{bmatrix}, \qquad X^T\boldsymbol{y}=\begin{bmatrix} 1 & 1 & 1 & 1 \\ 2 & 5 & 7 & 8 \end{bmatrix}\begin{bmatrix} 1 \\ 2 \\ 3 \\ 3 \end{bmatrix}=\begin{bmatrix} 9 \\ 57 \end{bmatrix}.$$

Solving the normal equations $\begin{bmatrix} 4 & 22 \\ 22 & 142 \end{bmatrix}\begin{bmatrix} a \\ b \end{bmatrix}=\begin{bmatrix} 9 \\ 57 \end{bmatrix}$, we find $\begin{bmatrix} a \\ b \end{bmatrix}=\begin{bmatrix} 2/7 \\ 5/14 \end{bmatrix}$.

Hence, the least squares line is $y=\frac{2}{7}+\frac{5}{14}x$.

The plot above shows our points together with this line. It does look like a very good fit!

Important comment. In what sense is this the line of "best fit"? By computing a least squares solution the way we do, we are minimizing the error $\boldsymbol{y}-X\begin{bmatrix} a \\ b \end{bmatrix}$. The components of that error are $y_i-(a+bx_i)$.

Hence, we see that we are minimizing the **residual sum of squares** $\mathrm{SS}_{\mathrm{res}}=\sum_i[y_i-(a+bx_i)]^2$.

Also see the discussion after the next example (where we swap the role of x and y) as well as the example at the beginning of next class (where we discuss making predictions and why minimizing $\mathrm{SS}_{\mathrm{res}}$ corresponds to minimizing the error of those predictions).

Armin Straub
straub@southalabama.edu

Example 46. (again) Determine the least squares line for the points $(2,1),(5,2),(7,3),(8,3)$.

Solution. Let's repeat the computation we did last class. This time, we let Sage do the actual work for us:

```
Sage] X = matrix([[1,2],[1,5],[1,7],[1,8]]); y = vector([1,2,3,3])
Sage] (X.transpose()*X).solve_right(X.transpose()*y)
```

$$\left(\frac{2}{7}, \frac{5}{14}\right)$$

Here are some intermediate steps to help see what's going on (and that it matches our earlier work):

```
Sage] X.transpose()*X
```

$$\begin{bmatrix} 4 & 22 \\ 22 & 142 \end{bmatrix}$$

```
Sage] X.transpose()*y
```

$$(9, 57)$$

Let's plot the least squares line $y = \frac{2}{7} + \frac{5}{14}x$ in Sage to marvel at the good fit!

```
Sage] points = [[2,1],[5,2],[7,3],[8,3]]
Sage] scatter_plot(points)
```

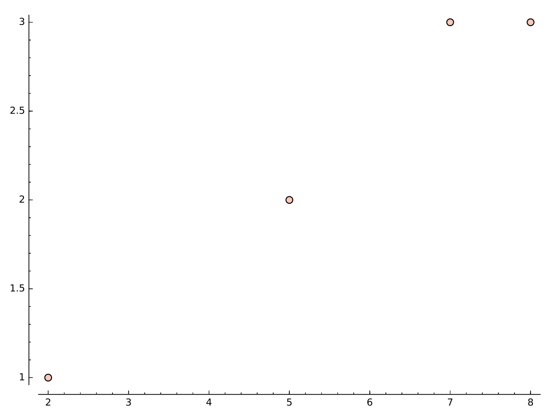


```
Sage] scatter_plot(points) + plot(2/7+5/14*x,1,9)
```

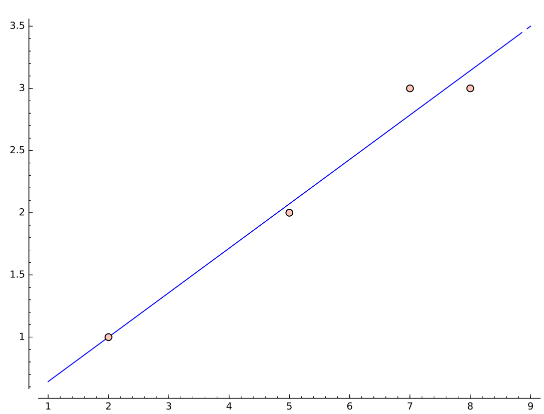

Comment. As mentioned earlier, the least squares line minimizes the (sum of squares of the) vertical offsets:

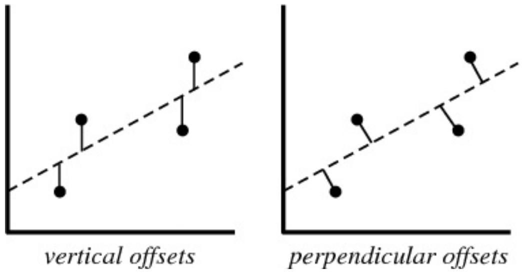

http://mathworld.wolfram.com/LeastSquaresFitting.html

Comment. We get a (slightly) different "best fit" line if we change the role of x and y! Can you explain that?

```
Sage] X = matrix([[1,1],[1,2],[1,3],[1,3]]); y = vector([2,5,7,8])
```

```
Sage] (X.transpose()*X).solve_right(X.transpose()*y)
```

$$\left(-\frac{7}{11}, \frac{30}{11}\right)$$

Note that $x = -\frac{7}{11} + \frac{30}{11}y$ is equivalent to $y = \frac{7}{30} + \frac{11}{30}x$.

```
Sage] scatter_plot([[2,1],[5,2],[7,3],[8,3]]) + plot(2/7+5/14*x,1,9) + plot(7/30+11/30*x,
      1,9,color='red')
```

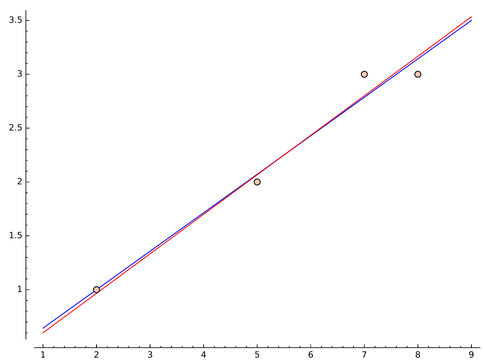

The explanation is that (see pictures at the beginning of this example) we are minimizing vertical offsets in one case and horizontal offsets in the other case.

In linear regression, the relationship between a dependent variable and one or more explanatory variables is modeled. If y is the dependent variable, with x the explanatory variable, then it is natural to minimize the error we make in "predicting y through x" (vertical offsets). See example at the beginning of next class!

Example 47. A car rental company wants to predict the annual maintenance cost y (in 100USD/year) of a car using the age x (in years) of that car (as an explanatory variable). Based on the observations $(x, y) = (2, 1), (5, 2), (7, 3), (8, 3)$, predict the cost for a 4.5 year old car (using linear regression).

Solution. Once we compute the regression line $y = a + bx$ (we already did that: $y = \frac{2}{7} + \frac{5}{14}x$), our prediction is $\frac{2}{7} + \frac{5}{14} \cdot 4.5 = \frac{53}{28} \approx 1.89$, that is, 189 USD/year.

Application: multiple linear regression

> *In statistics, **linear regression** is an approach for modeling the relationship between a scalar dependent variable and one or more explanatory variables.*
>
> *The case of one explanatory variable is called simple linear regression.*
>
> *For more than one explanatory variable, the process is called multiple linear regression.*
>
> http://en.wikipedia.org/wiki/Linear_regression

The experimental data might be of the form (x_i, y_i, z_i), where now the dependent variable z_i depends on two explanatory variables x_i, y_i (instead of just x_i).

Example 48. Set up a linear system to find values for the parameters a, b, c such that $z = a + bx + cy$ best fits some given points $(x_1, y_1, z_1), (x_2, y_2, z_2), \ldots$

Solution. The equations $a + bx_i + cy_i = z_i$ translate into the system:

$$\underbrace{\begin{bmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \\ 1 & x_3 & y_3 \\ \vdots & \vdots & \vdots \end{bmatrix}}_{\text{design matrix } A} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \underbrace{\begin{bmatrix} z_1 \\ z_2 \\ z_3 \\ \vdots \end{bmatrix}}_{\substack{\text{observation} \\ \text{vector } \boldsymbol{z}}}$$

Of course, this is usually inconsistent. To find the best possible a, b, c we compute a least squares solution by solving $A^T A \begin{bmatrix} a \\ b \\ c \end{bmatrix} = A^T \boldsymbol{z}$.

Application: Fitting data to other curves

We can also fit the experimental data (x_i, y_i) using other curves.

Example 49. Set up a linear system to find values for the parameters a, b, c that result in the quadratic curve $y = a + bx + cx^2$ that best fits some given points $(x_1, y_1), (x_2, y_2), \ldots$

Solution. $y_i \approx a + bx_i + cx_i^2$ with parameters a, b, c.

The equations $y_i = a + bx_i + cx_i^2$ in matrix form:

$$\underbrace{\begin{bmatrix} 1 & x_1 & x_1^2 \\ 1 & x_2 & x_2^2 \\ 1 & x_3 & x_3^2 \\ \vdots & \vdots & \vdots \end{bmatrix}}_{\text{design matrix} A} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \end{bmatrix}}_{\substack{\text{observation} \\ \text{vector } \boldsymbol{y}}}$$

Again, we determine values for a, b, c by computing a least squares solution to that system.

That is, we need to solve the system $A^T A \begin{bmatrix} a \\ b \\ c \end{bmatrix} = A^T \boldsymbol{y}$.

Example 50. (extra) Use Sage to find values for a, b, c that result in the quadratic curve $y = a + bx + cx^2$ that best fits the points $(0,1), (1,2), (2,3), (3,-4), (4,-7), (5,-12)$.

Solution. We first input the points:

```
Sage] points = [[0,1],[1,2],[2,3],[3,-4],[4,-7],[5,-12]]
```

We set up the system described in the previous example, then determine a least-squares solution.

```
Sage] X = matrix([[1,0,0],[1,1,1],[1,2,4],[1,3,9],[1,4,16],[1,5,25]])
```

```
Sage] y = vector([1,2,3,-4,-7,-12])
```

```
Sage] (X.transpose()*X).solve_right(X.transpose()*y)
```

$$\left(\frac{3}{2}, \frac{179}{140}, -\frac{23}{28}\right)$$

Hence, the best fitting quadratic curve is $y = \frac{3}{2} + \frac{179}{140} - \frac{23}{28}x^2$. Here's a plot:

```
Sage] scatter_plot(points) + plot(3/2+179/140*x-23/28*x^2,0,5,color='red')
```

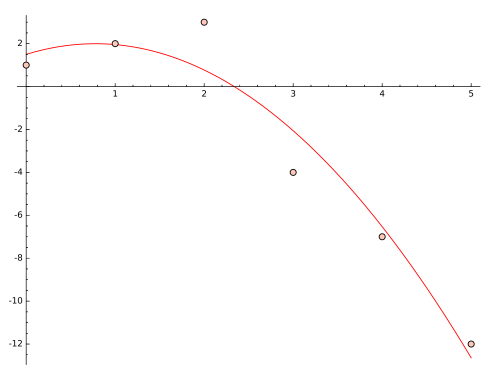

Advanced comment. If you are comfortable with Python, you can avoid typing out X and $\boldsymbol{y}$:
[The plot command above now won't work anymore because we are overwriting x with numbers.]

```
Sage] X = matrix([[1,x,x^2] for x,y in points])
```

```
Sage] y = vector([y for x,y in points])
```

More on orthogonality

Projection matrices

The **(orthogonal) projection** $\hat{\boldsymbol{b}}$ of a vector $\boldsymbol{b}$ onto a subspace W is the vector in W closest to $\boldsymbol{b}$.

We can compute $\hat{\boldsymbol{b}}$ as follows:

- Write $W = \operatorname{col}(A)$ for some matrix A.
- Then $\hat{\boldsymbol{b}} = A\hat{\boldsymbol{x}}$ where $\hat{\boldsymbol{x}}$ is a least squares solution to $A\boldsymbol{x} = \boldsymbol{b}$. (i.e. $\hat{\boldsymbol{x}}$ solves $A^T A\hat{\boldsymbol{x}} = A^T\boldsymbol{b}$)

Why? Why is $A\hat{\boldsymbol{x}}$ the projection of $\boldsymbol{b}$ onto $\operatorname{col}(A)$?

Because, for a least squares solution $\hat{\boldsymbol{x}}$, $A\hat{\boldsymbol{x}} - \boldsymbol{b}$ is as small as possible (and any element in $\operatorname{col}(A)$ is of the form $A\boldsymbol{x}$ for some $\boldsymbol{x}$).

Note. This is a recipe for computing any orthogonal projection! That's because every subspace W can be written as $\operatorname{col}(A)$ for some choice of the matrix A (take, for instance, A so that its columns are a basis for W).

Example 51. What is the orthogonal projection of $\begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}$ onto $\operatorname{span}\left\{\begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix}\right\}$?

Solution. In other words, what is the orthogonal projection of $\boldsymbol{b} = \begin{bmatrix} 2 \\ 0 \\ 11 \end{bmatrix}$ onto $\operatorname{col}(A)$ with $A = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}$.

In Example 44, we found that the system $A\boldsymbol{x} = \boldsymbol{b}$ has the least squares solution $\boldsymbol{x} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$.

The projection $\hat{\boldsymbol{b}}$ of $\boldsymbol{b}$ onto $\operatorname{col}(A)$ thus is $A\hat{\boldsymbol{x}} = \begin{bmatrix} 4 & 0 \\ 0 & 2 \\ 1 & 1 \end{bmatrix}\begin{bmatrix} 1 \\ 2 \end{bmatrix} = \begin{bmatrix} 4 \\ 4 \\ 3 \end{bmatrix}$.

Check. The error $\hat{\boldsymbol{b}} - \boldsymbol{b} = \begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix}$ needs to be orthogonal to $\operatorname{col}(A)$. Indeed: $\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix} \cdot \begin{bmatrix} 4 \\ 0 \\ 1 \end{bmatrix} = 0$ and $\begin{bmatrix} 2 \\ 4 \\ -8 \end{bmatrix} \cdot \begin{bmatrix} 0 \\ 2 \\ 1 \end{bmatrix} = 0$.

Example 52. (extra)

(a) What is the orthogonal projection of $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$ onto $\operatorname{span}\left\{\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}\right\}$?

(b) What is the orthogonal projection of $\begin{bmatrix} 1 \\ 2 \\ 2 \end{bmatrix}$ onto $\operatorname{span}\left\{\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}\right\}$?

Solution. (final answer only) The projections are $\left(\frac{11}{6}, \frac{1}{3}, \frac{7}{6}\right)^T$ and $\left(\frac{3}{2}, 0, \frac{3}{2}\right)^T$.

Armin Straub
straub@southalabama.edu