Before doing any math with a graph, ask: **What are the units?** A graph showing 120 packages might record a total, while a graph showing 120 packages per hour describes a rate. Those are different quantities, answer different questions, and need different interpretations.

A piecewise linear graph connects recorded data points with straight-line segments. It can be a useful way to summarize changing quantities, estimate values between observations, and make a complicated pattern easier to read. But a straight segment also makes a claim: it suggests that the quantity changed at a constant rate between the points. Whether that is a reasonable claim depends on the situation.

In this post, we'll define piecewise linear graphs, look at how to use linear interpolation, and discuss when this kind of model simplification is helpful—and when it can mislead.

What Is a Piecewise Linear Graph?

A **piecewise linear graph** is made of straight-line pieces. Each piece describes the relationship between two recorded points over a particular interval. The slope of each segment gives the rate of change over that interval.

Suppose a warehouse records the number of packages processed during a shift. At 9:00 a.m., the cumulative total is 40 packages, and at 10:00 a.m., it is 100 packages. Plotting those points and connecting them with a straight segment creates a piecewise linear model for the total between 9:00 and 10:00.

The segment's slope is:

$$ \frac{100\text{ packages} - 40\text{ packages}}{10:00\text{ a.m.} - 9:00\text{ a.m.}} = \frac{60\text{ packages}}{1\text{ hour}} = 60\text{ packages per hour.} $$

Notice the units: packages per hour. The recorded quantity is a total measured in packages; the slope is a rate measured in packages per hour. They are related, but they are not the same thing.

A piecewise formula

If we let t be the number of hours after 9:00 a.m. and $P(t)$ be the cumulative number of packages processed, the segment can be written as:

$$ P(t) = 40 + 60t, \quad 0 \leq t \leq 1. $$

The formula applies only on the stated interval. A piecewise model may use a different formula on the next interval because the rate may change. For example, the processing rate from 10:00 to 11:00 could be lower or higher than 60 packages per hour.

Start by Defining What the Graph Records

Before interpreting a graph, write one clear sentence describing what its vertical quantity means. For the warehouse example:

The graph records the cumulative number of packages processed since the shift began at 8:00 a.m.

That sentence clarifies both the quantity and its starting point. Without it, a reader might not know whether the graph shows the cumulative number of packages processed, the number processed during each hour, or the current processing rate.

The start time matters especially for cumulative totals. If a graph says $P(10)=100$, does 100 mean packages processed https://activecalculus.org/author/Running-Totals-Rates-of-Change since 8:00 a.m., since the start of the week, or since a counter was reset? A total is only interpretable when its counting period is clear.

Labels should make this information visible. A useful graph might label the horizontal axis “Time since shift began (hours)” and the vertical axis “Cumulative packages processed (packages).” Include the date or start time when it affects interpretation. Unlabeled axes leave readers guessing—and guesses about units can lead to bad decisions.

Totals and Rates Answer Different Questions

Ask what you want to know before choosing what to calculate. “How many packages have been processed?” asks for a **total**. “How quickly are packages being processed?” asks for a **rate**.

- **Total:** packages processed by a particular time, measured in packages.
- **Rate:** packages processed per unit of time, measured in packages per hour or packages per minute.
- **Change in total:** the difference between two cumulative totals, measured in packages.
- **Average rate of change:** change in the recorded quantity divided by elapsed time, measured in packages per hour.

For instance, if a cumulative total rises from 40 to 100 packages in one hour, the increase is 60 packages. The average rate over that hour is 60 packages per hour. Saying “the total is 60 packages per hour” mixes up a quantity with a rate.

This distinction appears in many workplace charts. An inventory dashboard might show units currently in stock (a total) or units sold per day (a rate). A finance chart might show an account balance (a total amount of money) or spending per month (a rate). Always check the vertical-axis label before interpreting the graph.

Intervals Matter for Average Rate of Change

An average rate is calculated over a specific interval. If $Q(t)$ is a recorded quantity, the average rate of change from time $t=a$ to time $t=b$ is:

$$ \fracQ(b)-Q(a)b-a. $$

The numerator is the change in quantity; the denominator is the elapsed time. The units of the rate follow from dividing the quantity’s units by the time units.

For example, suppose an inventory count is 500 units on Monday and 380 units on Friday. The change is -120 units over four days, so the average rate of change is:

$$ \frac-120\text units4\text days=-30\text units per day. $$

This does not mean the inventory fell by exactly 30 units every day. It means that the net change over that four-day interval averages out to a decrease of 30 units per day. The actual inventory could have risen on Tuesday, fallen sharply on Wednesday, and changed little on Thursday.

That is one reason to state the interval when reporting an average rate. A rate averaged over one hour can differ from a rate averaged over a full shift. The endpoints matter, too: choosing different start and end times can produce a different average.

Linear Interpolation: Estimating Between Recorded Points

Linear interpolation means estimating a value between two recorded points by assuming the quantity changed linearly between them. It is a practical way to **approximate between points** when measurements are sparse and a constant rate over the interval is plausible.

Suppose a delivery dashboard records 40 packages processed at 9:00 a.m. and 100 packages at 10:00 a.m. To estimate the cumulative total at 9:30 a.m., linear interpolation treats the rate as constant across that hour:

\[40\text packages + \left(60\text packages per hour\right)\left(0.5\text hours\right) =70\text packages. \]

The estimate is 70 packages. The units work out: packages per hour multiplied by hours gives packages, which can be added to the starting total.

But 70 is an estimate, not a known fact. It assumes the packages were processed at a steady average rate during the hour. If the team paused for a meeting and processed most packages afterward, the actual 9:30 total could be quite different.

Interpolation is not the same as extrapolation

Interpolation estimates *between* recorded points. Extrapolation estimates beyond the observed data. If the same line is extended to predict the 11:00 total, that prediction assumes the 9:00-to-10:00 rate will continue. That may be reasonable for a short period under stable conditions, but it is a stronger assumption than estimating within the measured interval.

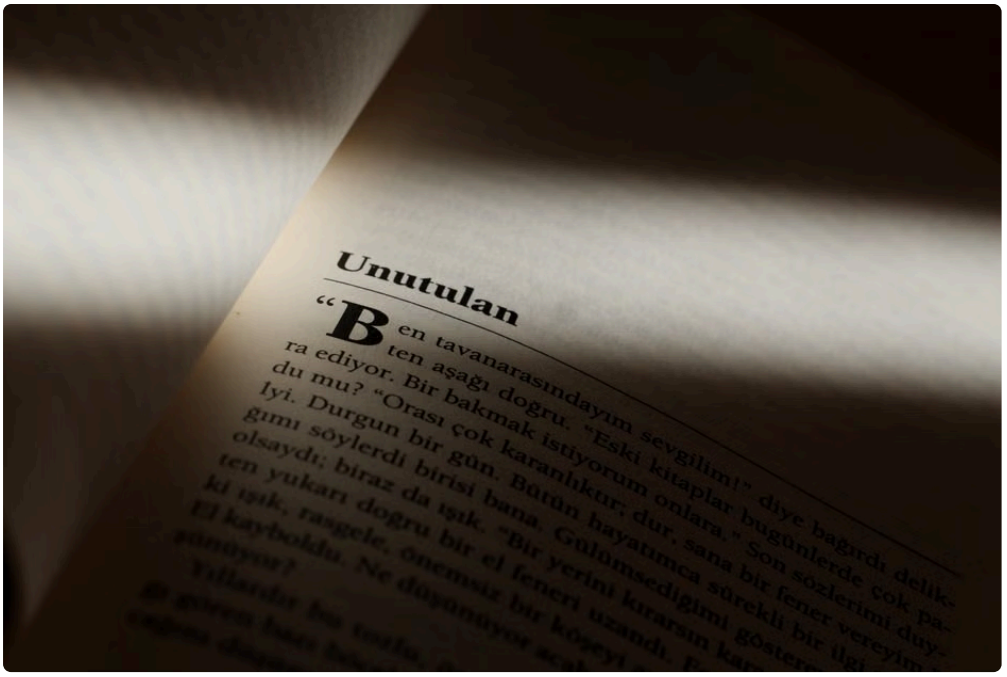
When Is a Piecewise Linear Model Okay?

A piecewise linear graph is often a good choice when the goal is a clear, useful summary rather than a perfect account of every moment. It can be appropriate when:

- Measurements are taken at distinct times and you need a readable picture of the trend.
- The quantity changes gradually enough that a constant average rate over each interval is a reasonable approximation.
- You are estimating between nearby observations and can clearly describe the estimate as approximate.
- The model's purpose is planning or communication, and the level of detail matches the decision being made.
- Different intervals may have different rates, so separate segments describe the pattern better than one long line.

For example, a shipping team might compare the cumulative number of orders packed at the beginning, middle, and end of each hour. Connecting those observations can help supervisors see when throughput appears to increase or slow down. The graph offers a useful overview, even though it does not reveal exactly when each order was packed.

This is **model simplification**: choosing a manageable representation of a more complicated process. A useful model does not need to reproduce every detail. It does need to make its assumptions visible and avoid suggesting more certainty than the data support.



When Straight-Line Connections Can Mislead

A straight segment visually represents change throughout the interval. If the points were recorded at 9:00 and 10:00, a line between them implies that the graph has a meaningful value at 9:30. That may be appropriate for an estimated cumulative total, but not for every kind of data.

Do not connect points automatically when:

- The quantity changes in sudden jumps, such as a balance changing only when a transaction posts.
- Measurements are snapshots and there is no reason to assume what happened between them.
- The process has known pauses, bursts, or thresholds that a straight segment hides.
- The interval is long enough for the rate to change substantially.
- The graph is being used for a decision that requires more accuracy than the available data provide.

For an account balance, a step graph may be more honest if deposits and withdrawals happen at specific times and the balance stays constant between transactions. Draw horizontal segments between events, then vertical jumps when a transaction occurs. That shape communicates “no change, then a sudden change” instead of implying that money entered or left continuously.

Likewise, if a counter updates only once per hour, connecting its readings with a diagonal line can imply that the counter recorded intermediate values. A line can still be used as an estimate, but label it as an approximation. If the real process changes in jumps, draw the jumps on purpose.

Changing Units Can Change the Story

Time conversions deserve care. If a rate is 60 packages per hour, then over 30 minutes the estimated increase is:

\[\left(60\frac\textpackages\texthour\right) \left(0.5\text hours\right)=30\text packages. \]

Using 30 as though it meant 30 hours would produce a wildly incorrect result. When mixing hours and minutes, convert one unit to the other before calculating, and write the units through the calculation so they can cancel correctly.

Units also help check whether a formula makes sense. A formula for a cumulative package total should produce packages. A formula for a processing rate should produce packages per hour. If the units do not match the question, the calculation is probably answering a different question.

A Practical Checklist for Reading or Making the Graph

1. **Name the recorded quantity.** Is it a total, a rate, a balance, or something else?
2. **State the start time or counting period.** What does the quantity count from?
3. **Label both axes with units.** Include time units and quantity units.
4. **Identify the interval.** Which two observations are being compared?
5. **Calculate the average rate carefully.** Use change in quantity divided by elapsed time.
6. **Decide how to show the in-between behavior.** Use straight segments for a reasonable linear approximation; use steps for event-driven jumps.
7. **Describe uncertainty.** Call interpolated values estimates, and avoid presenting extrapolations as measured facts.

Using MathJax and Runestone to Explore the Idea

MathJax is a tool that displays mathematical notation clearly on web pages. Expressions such as \(\fracQ(b)-Q(a)b-a\) are easier to read when fractions, variables, and units are rendered as math rather than left as plain-text shorthand. MathJax helps readers focus on the meaning of the formula instead of decoding its formatting.

Runestone is an interactive publishing platform used for online learning materials. An interactive activity can let learners adjust data points, compare segment slopes, and see how an estimate changes when the interval changes. That hands-on comparison reinforces a key idea: the average rate depends on the chosen interval, and the line between observations is a model of what may have happened—not necessarily a recording of every moment.

The Bottom Line

A piecewise linear graph connects recorded points with straight segments, using a separate segment for each interval. It is often useful for showing trends and making estimates between observations. Its main assumption is that the quantity changes at a constant rate within each segment.



Before trusting the picture, identify what is being recorded, clarify the start time, check the units, and name the interval. Keep totals distinct from rates. Use linear interpolation when a straight-line estimate is reasonable, and use a step graph when changes happen in jumps. A simple model can be valuable—but only when its assumptions match the question you are trying to answer.