


Describing Distributions

Time is money! Computers are fast, but so is human perception – and for software on the web, even small delays can lose you a lot of potential users. [Amazon once reported that delays as small as a tenth of a second cost them up to 1% of their overall revenue](#). There are software developers at every large software company who spend a significant amount of their time looking at graphs like the one below. (Your dear friend Mr. Kilian would know – he was one of them not too long ago 😊.)



This histogram describes the *response time* for a piece of software on the Internet. The response time is how long it takes for a software program to completely fulfill a user request. For example, when you click a link online, you might measure the response time between when you clicked the link and when the webpage finished loading.

Cloud software sits there all day, waiting for requests to come in and trying to fulfill them as fast as possible. There is natural variation in the response times for fulfilling those requests, which means response time isn't a number – it's a distribution! The histogram above shows the response time distribution for a real piece of software. The X-axis gives us response times, in milliseconds, grouped into 0.1 millisecond buckets. The Y axis tells us how often the response time for a request fell within each histogram bucket.

Your assignment: Write a few sentences describing the overall distribution graphed in the histogram above. Use the terms we covered in class today: for example, is this distribution unimodal or bimodal? Uniform or bell-shaped? Symmetric or skewed? Based on what you see in this graph, can you make any statements about how long this software generally takes to carry out its work?