



A picture with 8 colors



A picture with 256 colors



A picture with 16 million (24-bit) colors

As computer scientists, we often think about higher-level problems – designing algorithms, modeling problems with mathematical formulas, and analyzing theoretical bounds. But at the end of the day, everything we do on the computer is represented on disk and in memory as a series of 1's and 0's. It's important to represent data compactly and efficiently, and different forms of encoding and compression help us do that.

You may know how many megapixels your cell phone camera has, but do you know what that means? (Hopefully you do now! 😊) Consider the pictures at the left. As we increase the number of colors per pixel, the picture becomes clearer and crisper.

Pictures, video, music, text – each kind of data is encoded and compressed in different formats. Ideally, compression should be *lossless* - that is, the original data can be accurately reconstructed from the compressed data. But it's not always realistic or necessary. For example, can you tell which of the two pictures below is a compressed JPEG instead of a bitmap? The JPEG is *lossy*, meaning it loses some data, but it only takes up $\frac{1}{4}$ the space on disk. (It's the one on the right.)

Related Columbia Classes

Core classes:

COMS W3203 Discrete Mathematics

Advanced classes:

COMS W4160 Computer Graphics

COMS W4165 Computational Techniques in Pixel Processing

COMS W4167 Computer Animation



A picture with 16 million (24-bit) colors