Probabilistic inference provides a language for describing how organisms may learn from and adapt to their environment. The computations needed to implement probabilistic inference often require specific representations, akin to having the suitable data structures for implementing certain algorithms in computer programming. Yet it is unclear how such representations can be instantiated in the stochastic, parallel-running biochemical machinery found in cells (such as single-celled organisms). I will discuss how representations needed to support inference in Markov models can be embodied in cellular circuits. I will argue that the logic of protein production and degradation constrains the computation we set out to implement. This process, in which an abstract computation is shaped by its biochemical realization, strikes a compromise between "rationalistic" information-processing perspectives and alternative approaches that emphasize embodiment.