



Computer vision enables all sorts of new, powerful embedded systems, like self-driving cars and autonomous robots. Today's powerful hardware for embedded vision allows for great performance, but how do you ensure your solution makes efficient use of the hardware, for maximum performance?