July 2022

Power Consumption Predictions for Computing Devices During Design Phase

Ruben Rodriguez Buchillon

Udaya Kiran Ammu

Power Consumption Predictions for Computing Devices During Design Phase

ABSTRACT

This disclosure describes techniques to predict power consumption of a computing device under design. Per techniques of this disclosure, a machine learning model is trained based on parameterized hardware attributes of existing devices. A training dataset is obtained based on device parameters and power consumption of existing computing devices. The generated training dataset is utilized to train a machine learning (ML) model which is tested using a variety of hardware combinations. The ML model is verified using a test scenario and computing device configurations that were excluded from the training dataset. Upon successful verification, the ML model can be used for estimating power consumption of computing devices under design based on their component combinations.

KEYWORDS

- Device design
- Laptop design
- Component selection
- Component combination
- Power consumption
- Energy consumption
- Battery size
- Battery life

BACKGROUND

Computing devices, including portable devices such as laptops, can be built from a large selection of possible hardware components in a variety of combinations. During the design phase for a device, domain experts estimate power consumption of the yet to be built device based on heuristics, component datasheets, and prior design power measurements where available. However, this approach can fail to capture non-linear relationships in the interplay between components. Correct prediction of power consumption of a device design is key to correctly sizing the battery. Incorrect battery sizing in either direction can negatively affect the consumer experience since battery undersizing can lead to poor battery life and battery oversizing can lead to a higher product cost.

DESCRIPTION

This disclosure describes techniques to predict the power consumption of a computing device during the design phase. Per techniques of this disclosure, a machine learning model is trained based on parameterized hardware attributes of existing devices. The trained machine learning model is subsequently utilized to predict the power consumption of a newly designed computing device with a novel (not previously used) combination of hardware components.

Parametrization of the hardware attributes enables the machine learning model to learn the relationships between parameters of components, rather than rely on individual component specifications, as provided by their manufacturer. Parametrizing the device hardware attributes during a training phase of the machine learning model additionally enables the model to share training data points and enables characterization of a class of components rather than individual component versions.

The use of a machine learning model additionally enables characterization of relationships between different components that can influence power consumption, e.g., between a display panel and a processor. For example, a higher resolution display panel can lead to a higher current consumption in order to drive the pixel logic. However, it can also cause higher processor activity to generate the larger canvas as well as the need for increased content modifications such as upsampling/downsampling.

The relationships between different hardware components can be both linear and nonlinear. The trained machine learning model can provide accurate predictions on previously unseen/unmeasured components, based on the component parameters and their datasheet information. Additionally, parametrizing the display panel enables the model to learn underlying structures that all high-resolution panels might share (e.g., higher GPU activity), or even the relationship between increasing a display resolution of a device and its collateral effects on other components.

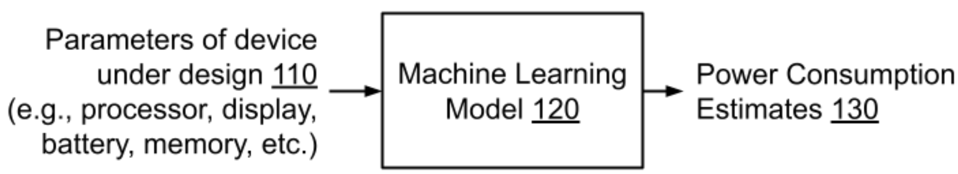

Fig. 1: A trained machine learning model provides power consumption estimates

Fig. 1 is a schematic of an example of a machine learning model utilized to provide power consumption estimates, per techniques of this disclosure. As depicted in Fig. 1, during a design phase, device parameters (110) of a proposed computing device configuration are provided to a trained machine learning model (120). The trained machine learning model provides power consumption estimates (130) for various modes of operation of the proposed device. The power consumption estimates can be utilized by designers to select various components for the device under design.

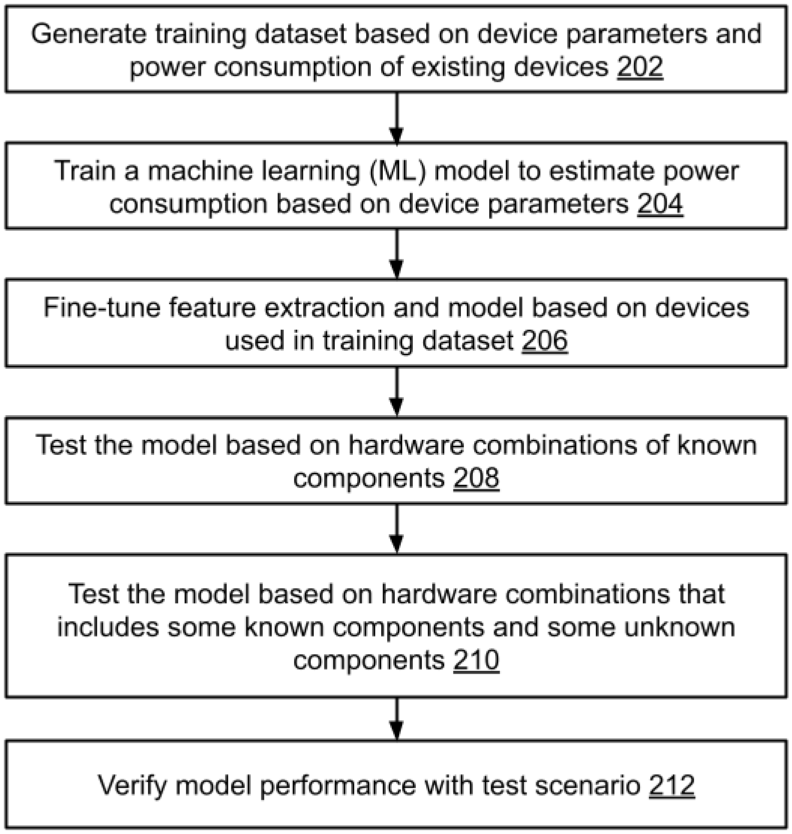

Fig. 2: Training and utilizing a machine learning model for computing device design

Fig. 2 depicts an example method to train and utilize a machine learning model to predict power consumption of a computing device under design, per techniques of this disclosure. A training dataset is generated (202) based on device parameters and power consumption of existing computing devices. In some cases, additional power measurements can be made using standard test cases that align well with real life use of devices.

The generated training dataset is utilized to train (204) a machine learning (ML) model. A neural network architecture may be utilized for the machine learning model. The training can utilize a hand-tuned feature extraction pipeline. In order to train the ML model, feature extraction is performed by determining, for each component in a computing device, details such as parameters to be extracted, datasheet power number estimates to be incorporated, power measurements, procedure(s) to handle missing data/components, encoding of semantic (e.g., text based) parameters, etc.

After the machine learning model is trained, multi-stage testing can be performed. The model is tested using hardware combinations (test points from devices utilized in training) that were provided to the model during training to verify whether the ML model is able to reconstruct relationships that were provided in training. This stage is utilized to fine-tune (206) the feature extraction and basic architecture of the ML model. The model is then tested (208) against hardware combinations where each component was seen explicitly in training, but that were not combined in any single computing device.

The model is further tested (210) using hardware combinations where some of the components were not in the configuration of any of the devices in the training data. This stage is used to verify whether the parametrization of the components provides sufficient data and whether device power consumption can be accurately predicted by the ML model based on a combination of component parameters.

The ML model is then verified (212) using a test scenario and a computing device that was not utilized for training the ML model. The test scenario can be an “IDLE ON” scenario where the computing device is turned on, the display is turned on, the user is logged in, and idle, and has invoked an application, e.g., opened a new browser tab. This test scenario is utilized since it is a stable scenario in which the device consumes sufficient power to enable power measurement while minimizing measurement bias, and whereby the device has all its components turned on, even if in idle state.

Verification of the ML model is considered successful if the predicted power consumption from the ML model matches the actual power consumption measured for the computing device. Upon successful verification, the ML model can be used for estimating power consumption of computing devices under design that use novel combinations of components.

Techniques of this disclosure can enable reliably and accurate prediction of energy (power) consumption of a computing device under design based on the combination of components. The early prediction and risk detection enables computing device manufacturers to make better hardware selections early in the design phase, thereby reducing cost, test burden, and uncertainty.

CONCLUSION

This disclosure describes techniques to predict power consumption of a computing device under design. Per techniques of this disclosure, a machine learning model is trained based on parameterized hardware attributes of existing devices. A training dataset is obtained based on device parameters and power consumption of existing computing devices. The generated training dataset is utilized to train a machine learning (ML) model which is tested using a variety of hardware combinations. The ML model is verified using a test scenario and computing device configurations that were excluded from the training dataset. Upon successful verification, the ML model can be used for estimating power consumption of computing devices under design based on their component combinations.

REFERENCES

1. Tae-Young Kim, Sung-Bae Cho, "Predicting residential energy consumption using CNN-LSTM neural networks." Energy, Volume 182, 2019, Pages 72-81, ISSN 0360-5442, https://doi.org/10.1016/j.energy.2019.05.230, accessed on 21 June 2022