Learning to Estimate Nutrition Facts from Food Descriptions

Hadi Amiri, Andrew L. Beam, Isaac S. Kohane
Harvard University
hadi@hms.harvard.edu
http://scholar.harvard.edu/hadi/

HARVARD MEDICAL SCHOOL | DEPARTMENT OF Biomedical Informatics

Disclosure

- I have nothing to disclose.

Food Computing

Computational models to acquire and analyze heterogeneous food data from disparate sources for:

- Recognition
- Retrieval
- Recommendation
- Monitoring

Food Computing

- Food choices can help preventing chronic diseases such as heart disease, diabetes, stroke, and certain cancers.
- Concerned with healthy eating & monitoring the nutrients that we consume.
- The availability of large-scale food datasets can transform the way individuals consume food.

Eppinger, A., et al. 2018. Serum Marker of Glyphosate Exposure Associated with Changes in Oral and Gut Microbiome Composition. AMIA'18.

Food Computing – Major Goal

- The ability to promote healthy eating through monitoring nutrients that we consume.
  - Need to either
    - have prior info about nutrition facts of foods, or
    - estimate such info.

Food Computing – Recent Efforts

- Large-scale food datasets such as USDA
- Computational models to *match* given meal descriptions with foods that exists in databases like USDA.
- Effective for foods that exists in databases.

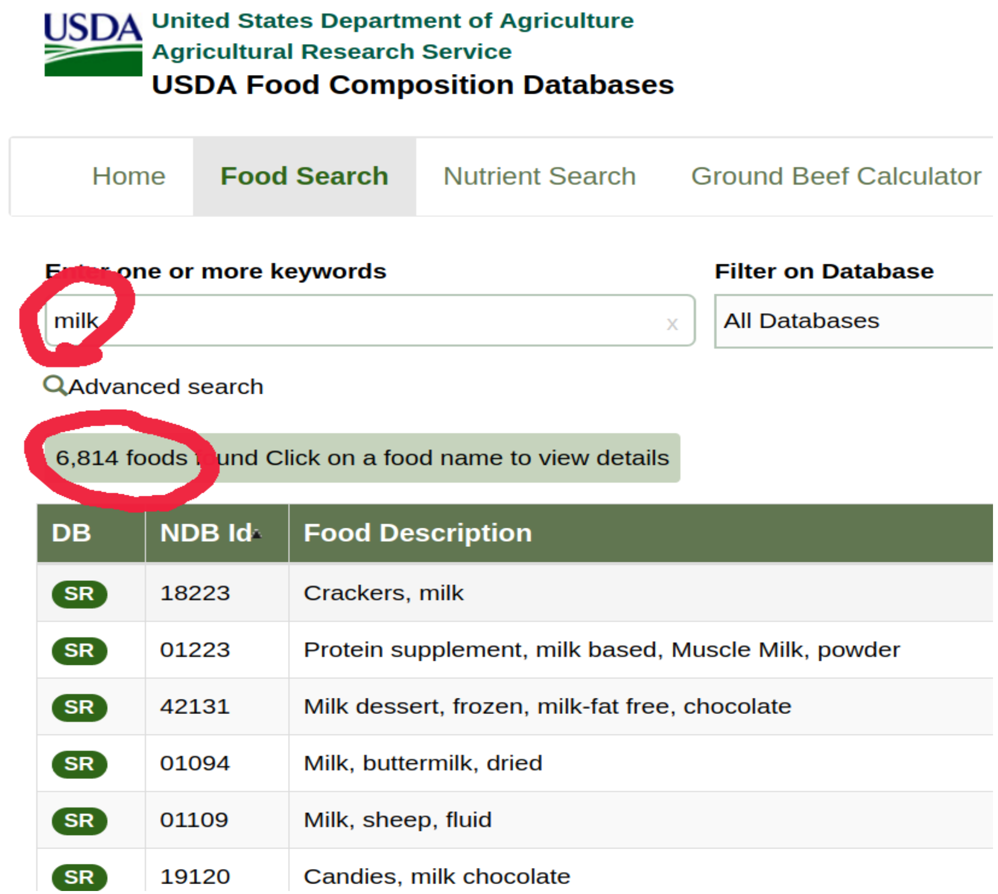

Korpusik, M., et al. 2017. Semantic mapping of natural language input to database entries via convolutional neural networks. ICASSP'17.

Food Computing – Challenges

- Avg. number of **new** foods per year: ~ 20K
  - USDA data is updated on a **yearly basis**.

- Previous approaches lack the ability to deal with **new foods** that don't exist in databases.

- Useful info such as **ingredients** have not been utilized to learn nutrition facts of food.

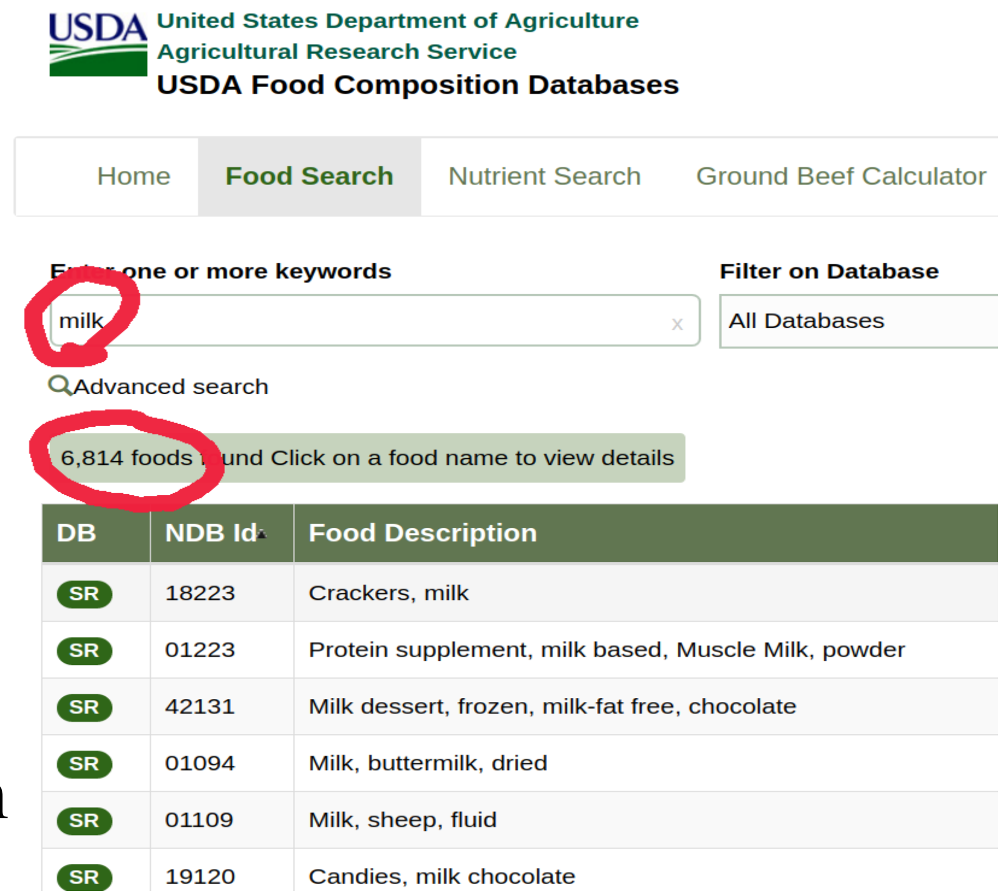

Study Aims

- To develop effective computational models to ***estimate*** nutrition facts of ***any*** given food, and

- To investigate if computational modeling of food ***ingredients*** can help better estimation of ***nutrition facts***.

USDA Dataset

- USDA Dataset
  - 237K food items
  - 40 nutrition fact types
    - carbohydrate, energy (calories), lipid, etc.
  - 100K ingredient types
    - sugar, salt, oil, etc.

USDA Data: http://www.ars.usda.gov/nutrientdata

USDA Dataset

- USDA Dataset
  - 237K food items
  - 40 nutrition fact types
  - 100K ingredient types

- Voluntarily supplied by food industry organizations to USDA.
- USDA standardizes the reported nutrition facts

USDA Dataset

- USDA Dataset
  - 237K food items
  - 40 nutrition fact types
  - 100K ingredient types

- Missing data
  - some organizations provide partial information about their products.
- Yearly updates
  - the dataset is updated on a yearly basis missing many new foods.

USDA Dataset

- USDA Dataset
  - 237K food items
  - 40 nutrition fact types
  - 100K ingredient types

- Power law distribution in nutrition facts: `c × exp(-0.14x)`
  - some nutrition facts match with only a small number of foods.

Method

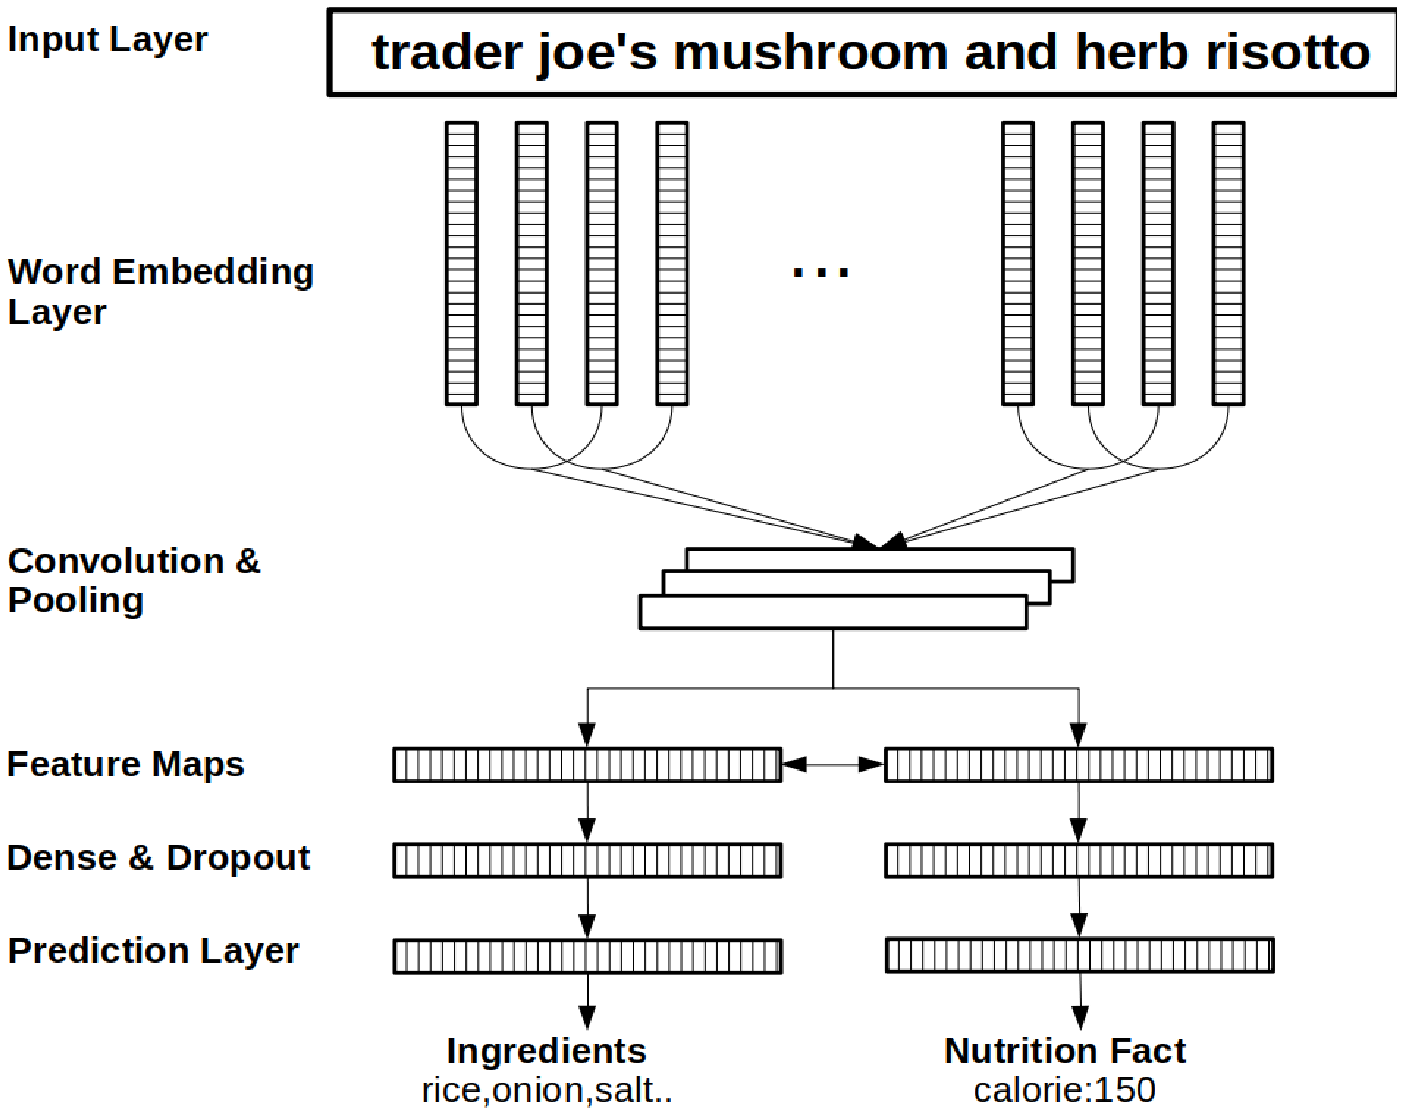

Multi-task learning framework (M-CNN)

- Joint learning of nutrition facts and ingredients given food descriptions.
- The shared layers were used to exploit commonalities and differences across tasks.
- Nutrition facts were normalized scalars.
- Framework was trained and tested on each nutrition fact separately.
- Ingredients of each food were represented by a vector of 0/1 vectors.
- The network (CNN) was trained by minimizing the MSE losses:

$$\mathcal{L}(I) = \mathcal{L}_{nutrition}(I) + \alpha \times \beta \times \mathcal{L}_{ingredient}(I)$$

α : controls ingredient contributions

β : establishes a common scale for losses (set to loss ratio computed at first iteration)

Method

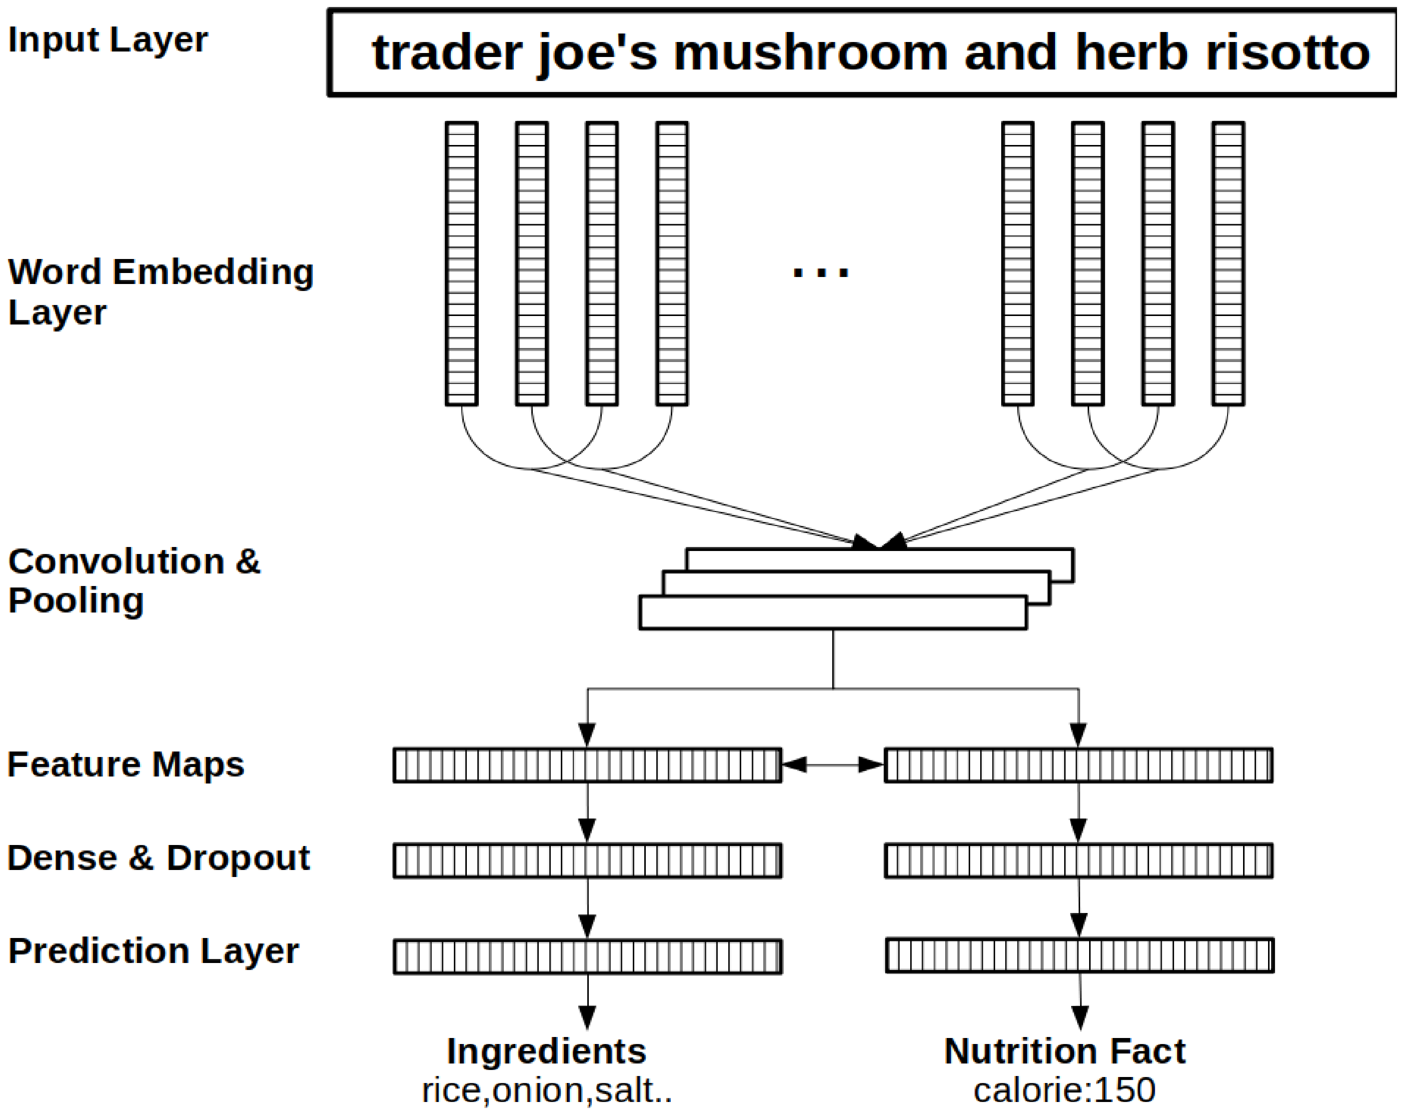

Multi-task learning framework (M-CNN)

- Learns semantic relations between *food items* and *ingredients*
  - learning that "roasted" foods should have "oil" as their ingredients
- Learns semantic relations between *food items* and *nutrition facts*
  - learning that "rice" has high "calories."

$$\mathcal{L}(I) = \mathcal{L}_{nutrition}(I) + \alpha \times \beta \times \mathcal{L}_{ingredient}(I)$$

α : controls ingredient contributions

β : establishes a common scale for losses (set to loss ratio computed at first iteration)

Method

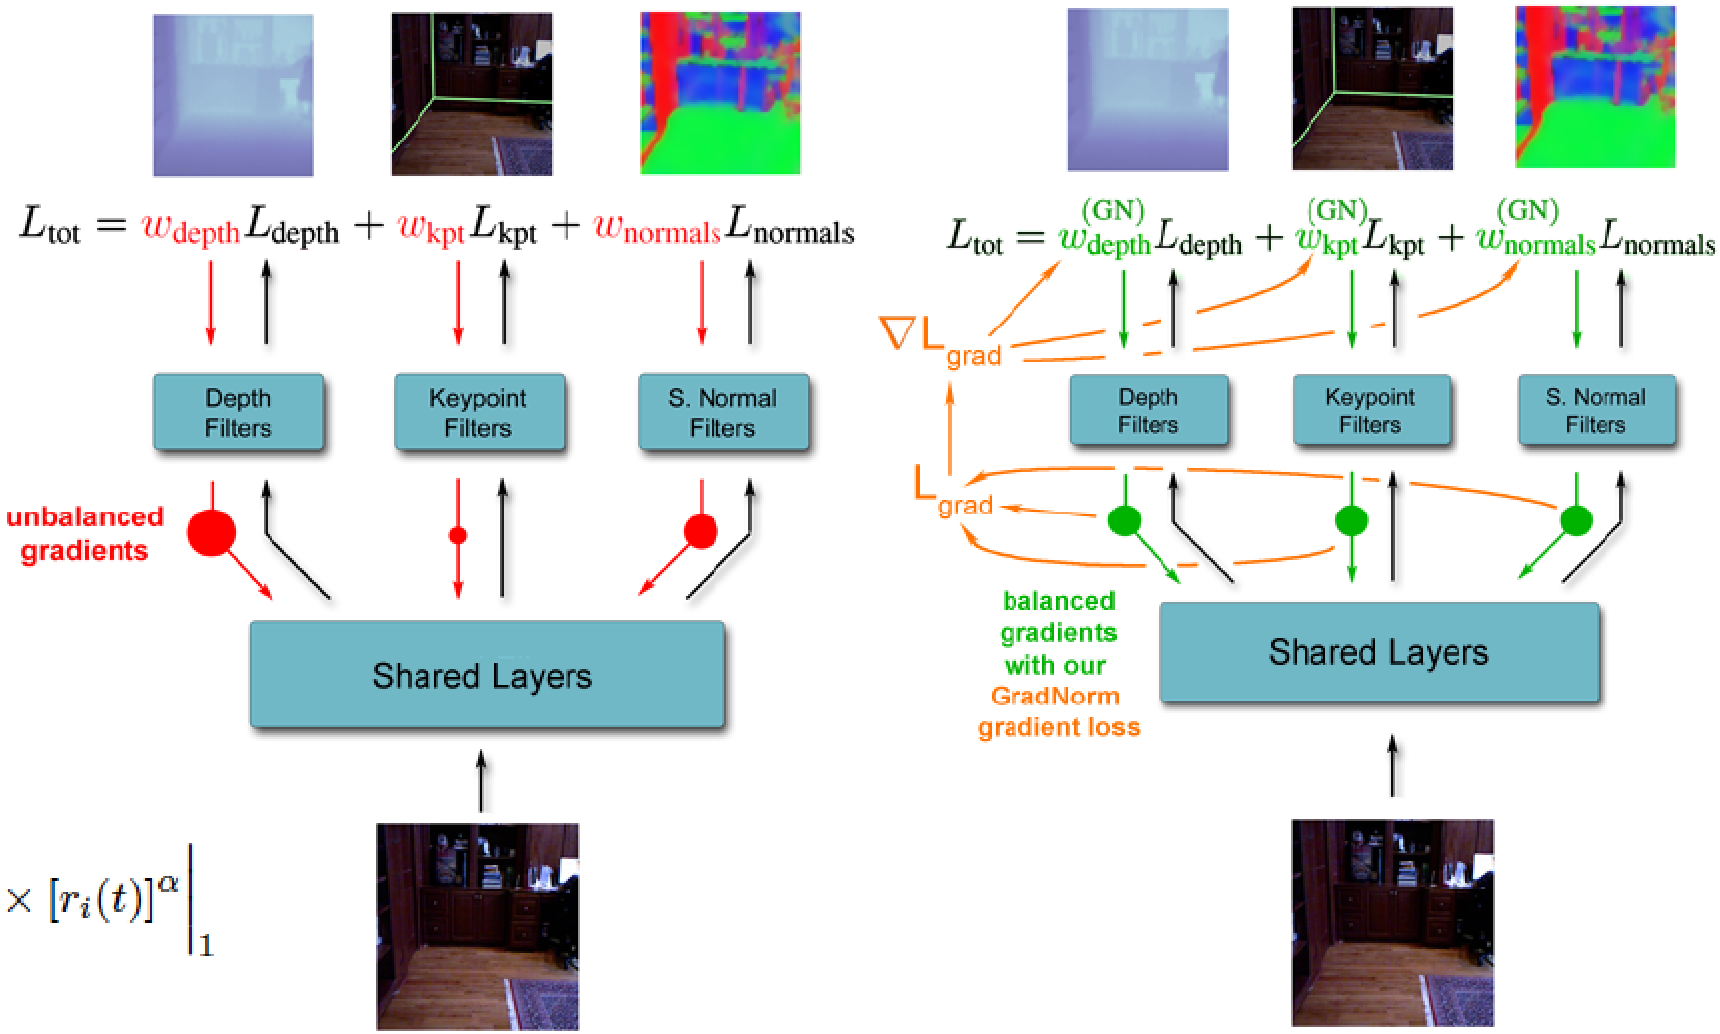

$$G_W^{(i)}(t) \mapsto \overline{G}_W(t) \times [r_i(t)]^\alpha,$$

$$L_{grad}(t; w_i(t)) = \sum_i \left| G_W^{(i)}(t) - \overline{G}_W(t) \times [r_i(t)]^\alpha \right|_1$$

Not been utilized in our work. Hopefully in future.

Figure 1. **Gradient Normalization.** Imbalanced gradient norms (left) result in suboptimal training within a multitask network. We implement GradNorm through computing a novel gradient loss L_{grad} (right) which detects such imbalances in gradient norms amongst tasks and tunes the weights w_i in the loss function to compensate. We illustrate here a simplified case where such balancing results in equalized gradient norms, but in general there may be tasks that require relatively high or low gradient magnitudes for optimal training (discussed further in Section 3).

Chen, Zhao, et al. Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. ICML'18.

Results

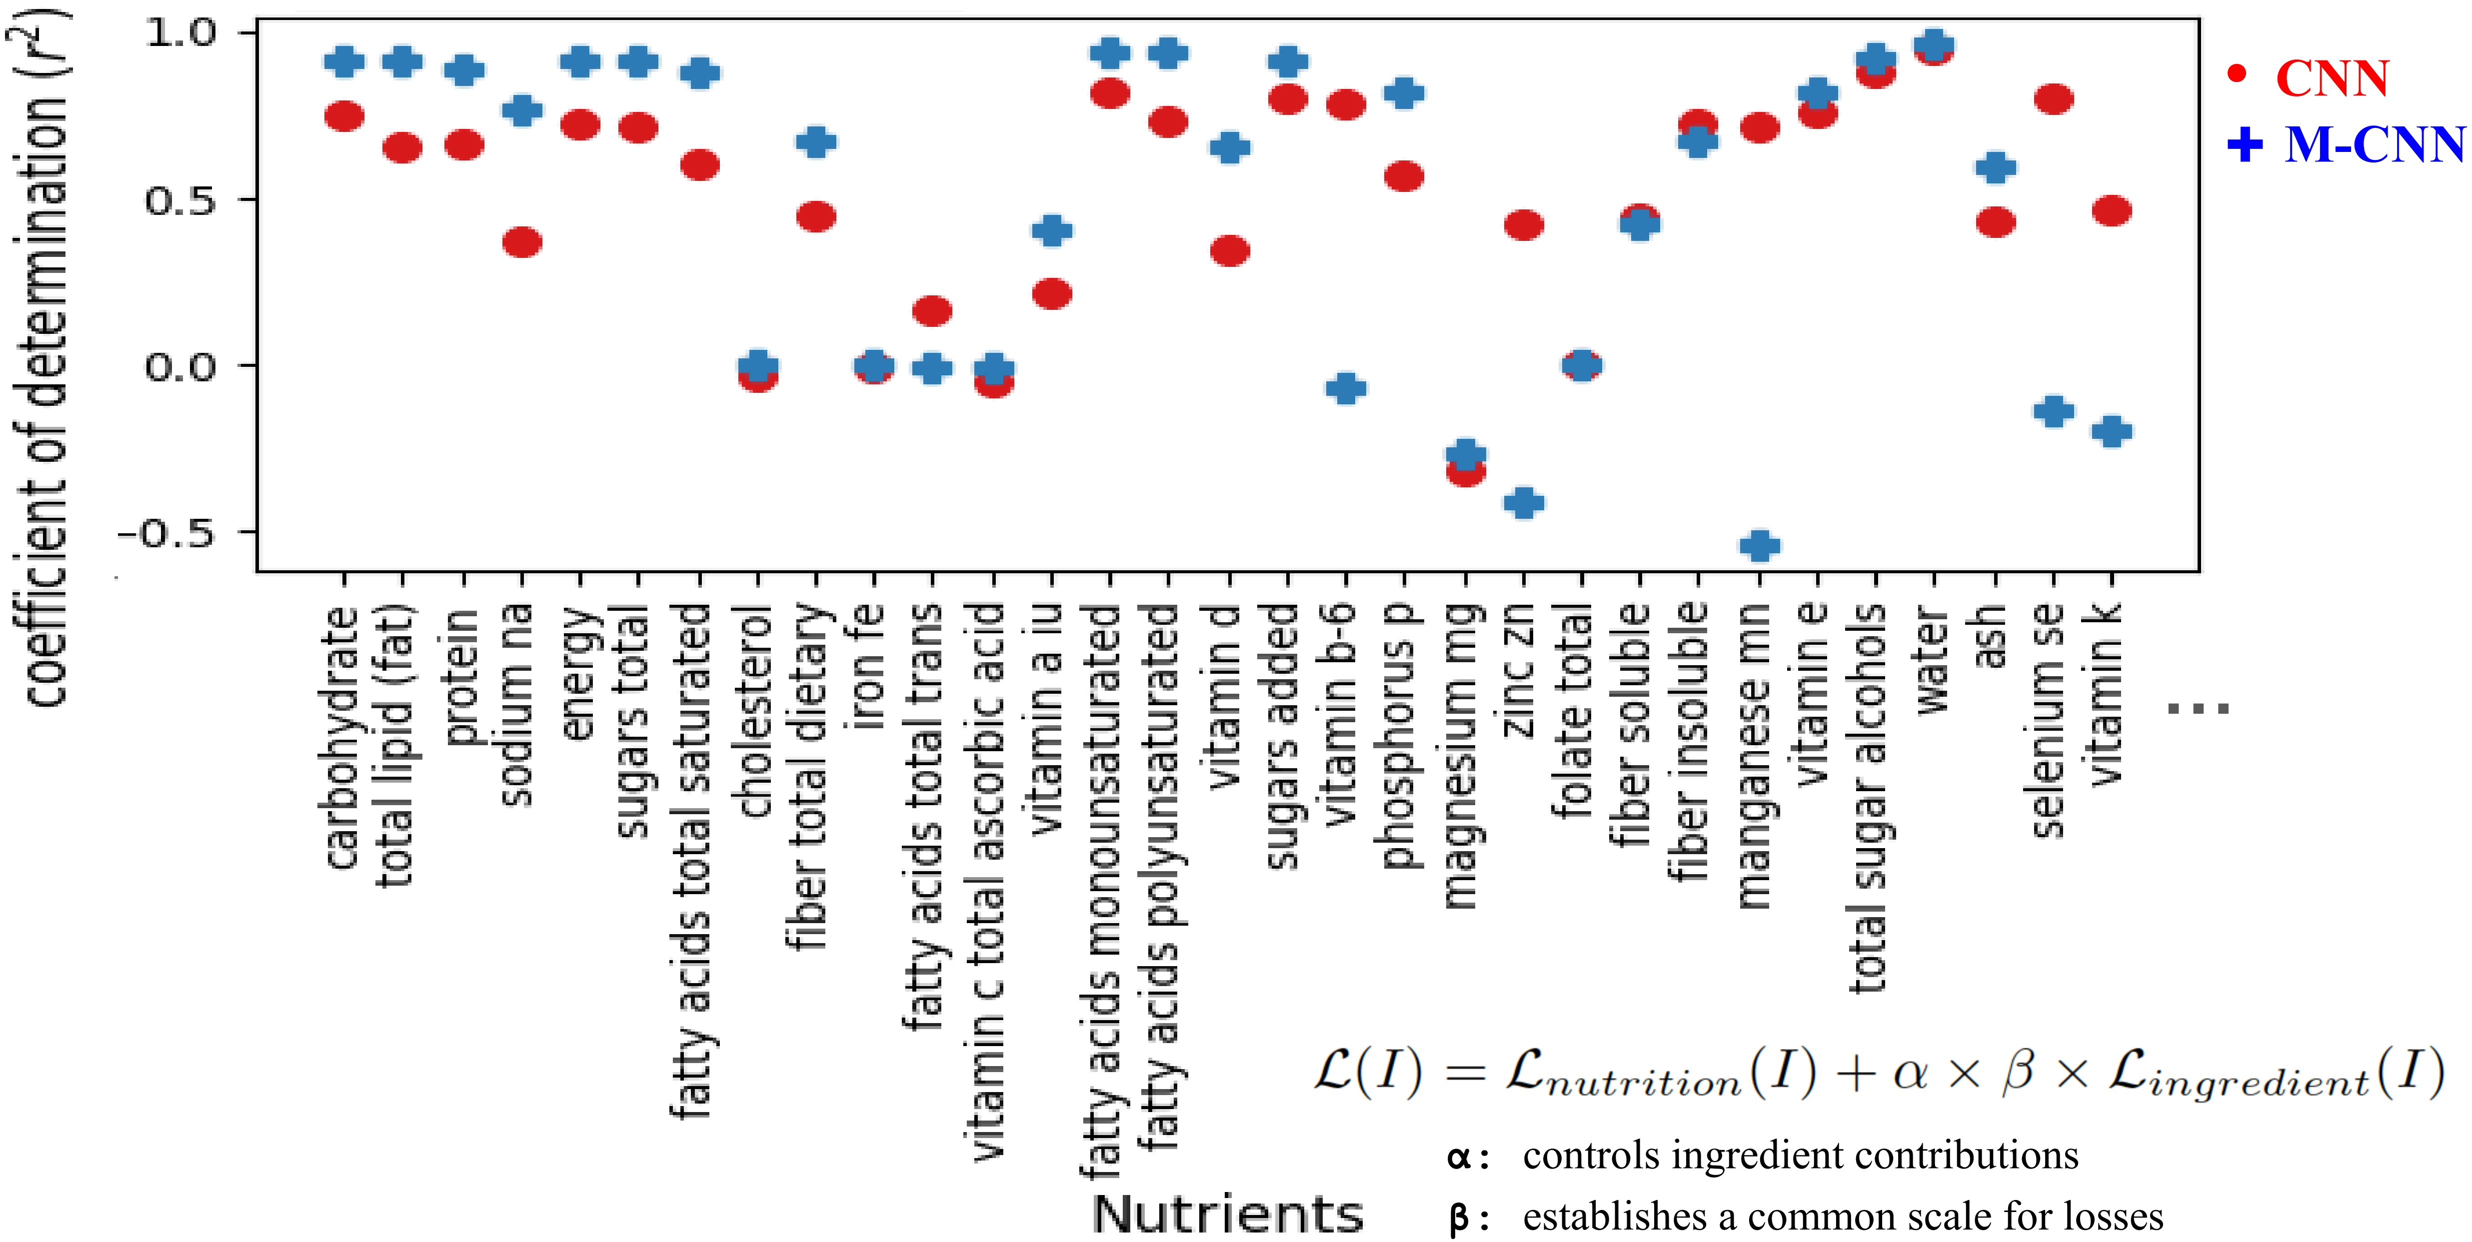

$$\mathcal{L}(I) = \mathcal{L}_{nutrition}(I) + \alpha \times \beta \times \mathcal{L}_{ingredient}(I)$$

α : controls ingredient contributions

β : establishes a common scale for losses

Results

$$\mathcal{L}(I) = \mathcal{L}_{nutrition}(I) + \alpha \times \beta \times \mathcal{L}_{ingredient}(I)$$

α : controls ingredient contributions

β : establishes a common scale for losses (set to loss ratio computed at first iteration)

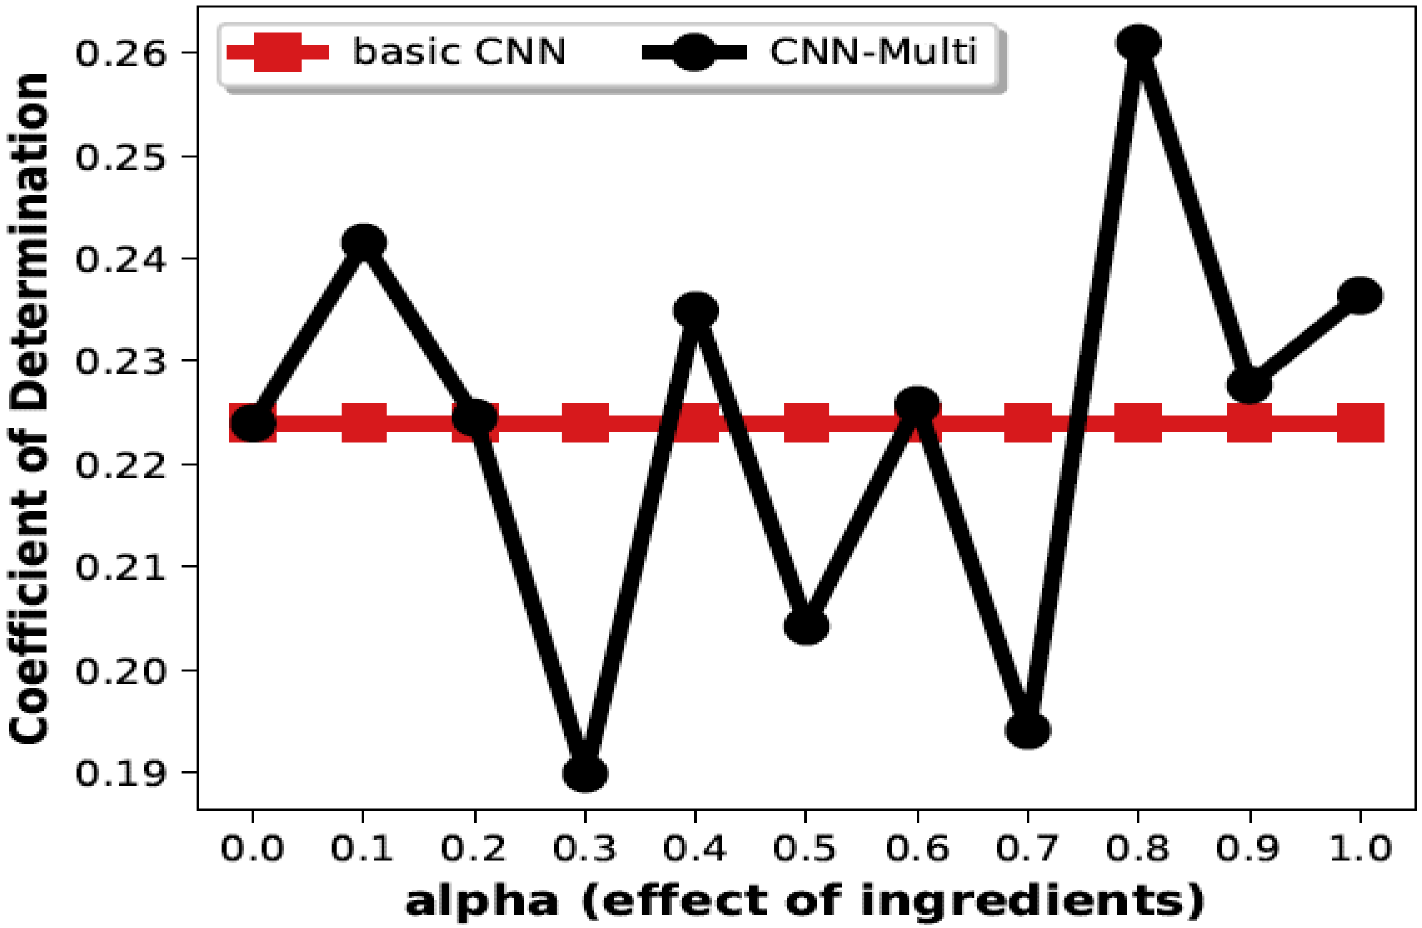

- α = 0 led to R2 of 0.224, while there existed other α values, i.e. α={.1, .4, .8, .9, 1}, that further improved the performance.
- We attribute this improvement to our model's ability in utilizing semantic relations between food items, and their ingredients and nutrition facts.

Conclusions

- Developed an effective regressor to accurately estimate nutrition facts of foods from their short descriptions.
- Highlighted the importance of learning ingredients for accurate estimation of nutrition facts.
- Our research can be used in diet monitoring applications, with significant public health impact.

Future Work

- Explore other factors such as food *quantity* and *type*.
- Ingredients often have a hierarchical form which could be utilized to create a better semantic space for ingredients.

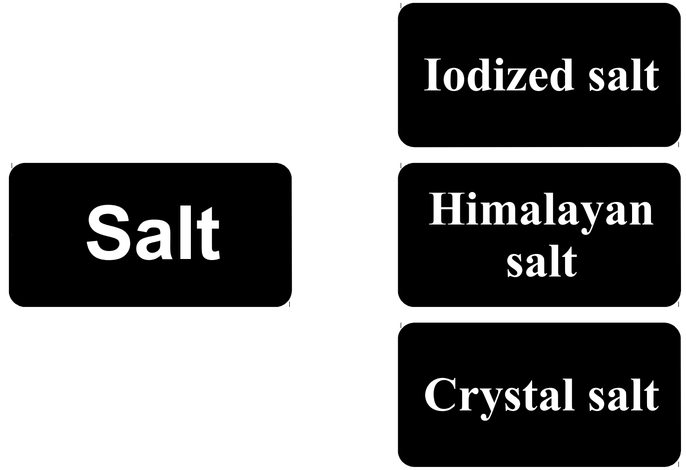

- Our learning framework is trained on each nutrition fact separately; joint learning of these facts might create stronger regressors.
- Extension to other applications such as
  - Food-drug interaction

Thank You!

Contact

Hadi Amiri

hadi@hms.harvard.edu

http://scholar.harvard.edu/hadi/

HARVARD MEDICAL SCHOOL

DEPARTMENT OF Biomedical Informatics