Multi-Task Learning for Calorie Prediction on a Novel Large-Scale Recipe Dataset Enriched with Nutritional Information

Robin Ruede, Verena Heusser, Lukas Frank, Alina Roitberg, Monica Haurilet, Rainer Stiefelhagen

Karlsruher Institut für Technologie

cv:hci

Nutrition Facts (per 100 g)		
	Pred	True
Calories	183 kcal	198 kcal
Fat	9 g	9 g
Carb	17 g	24 g
Protein	7 g	4 g
Ingredients (pred): Flour, Butter, Milk		
Ingredients (true): Eggs, Flour, Vanilla Sugar		

In a Nutshell

- A method for collecting nutritional information of recipes by aggregating semi-structured ingredient data

- *pic2kcal* benchmark: The largest dataset of images with calories captured in the wild

- Multi-task prediction of nutritional values and ingredients from pictures

Dataset

- Most existing datasets are domain specific and small
- Our dataset comprises a diverse set of recipe types and cuisines

Dataset Creation Pipeline

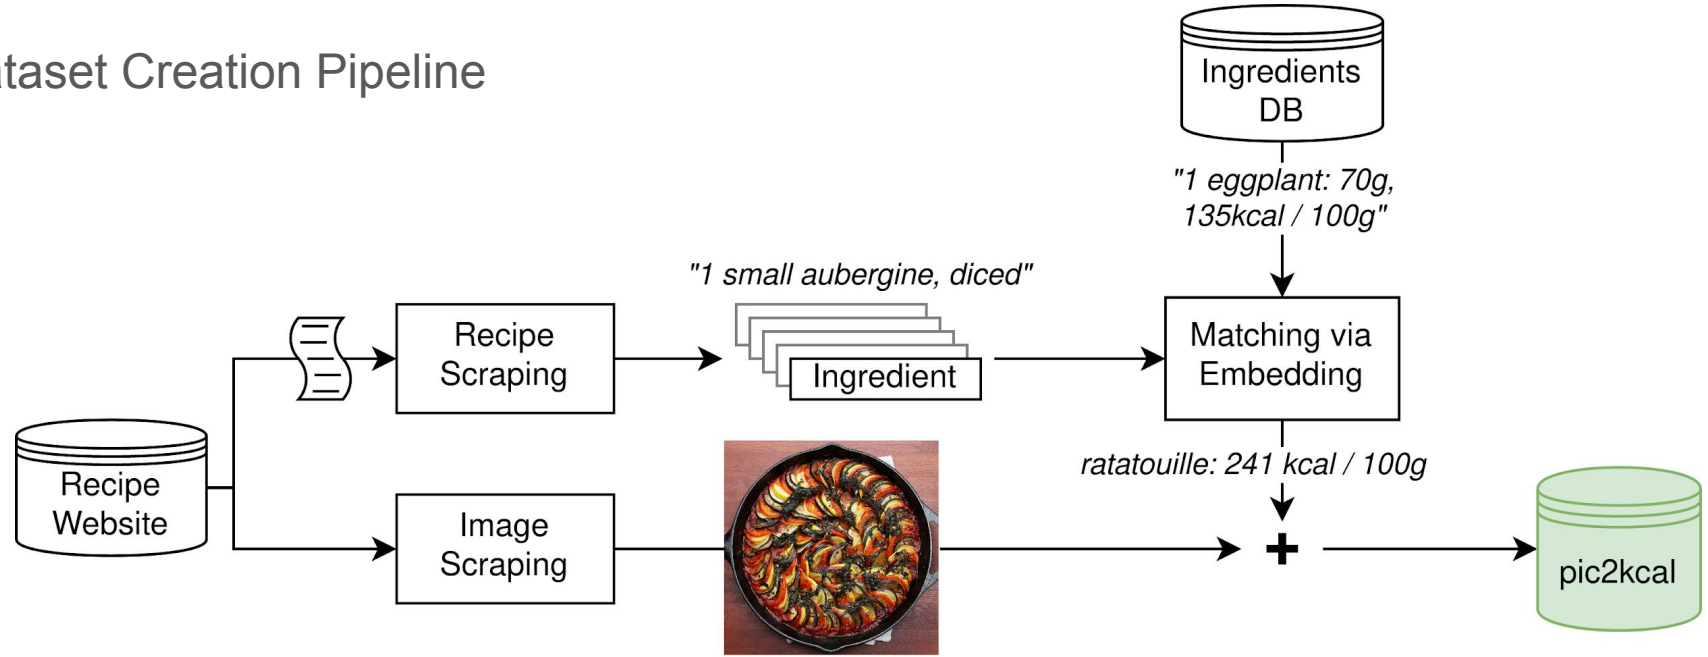

Multi-Task Learning for Calorie Prediction on a Novel Large-Scale Recipe Dataset Enriched with Nutritional Information, Ruede et al.

Dataset Statistics

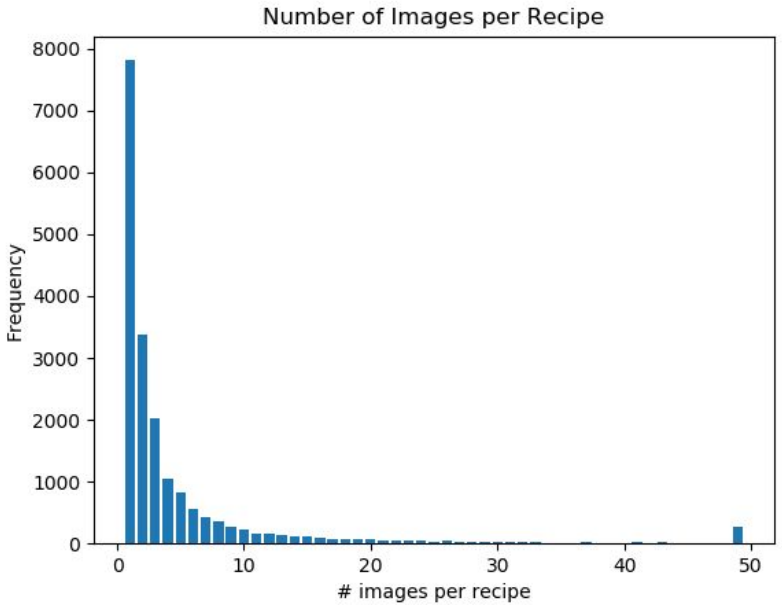

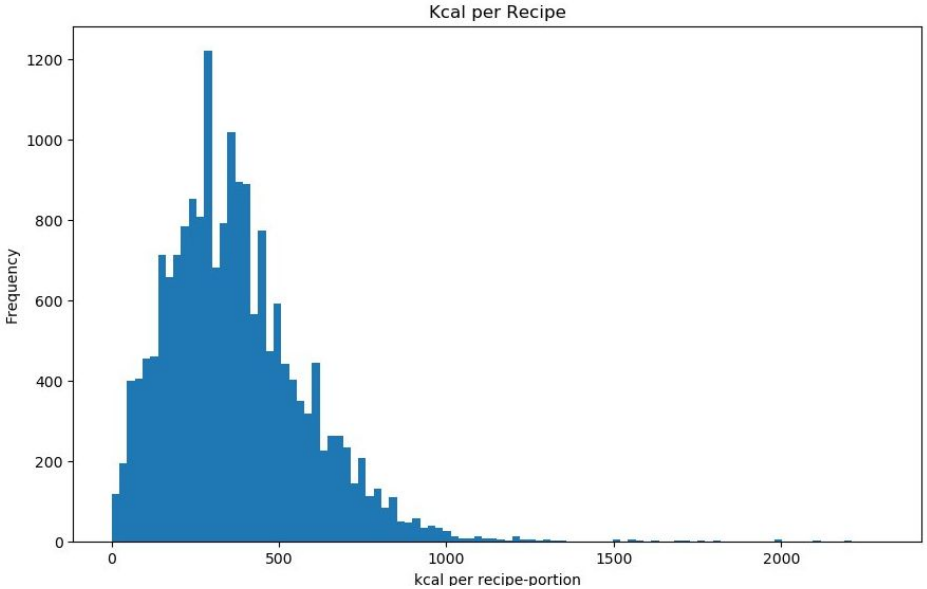

Multi-Task Learning for Calorie Prediction on a Novel Large-Scale Recipe Dataset Enriched with Nutritional Information, Ruede et al.

Dataset Comparison

By number of images with information calorie density per X

Dataset	Property	Per portion	Per 100g	Per recipe
Recipe1M+	Mean [kcal]	N/A	219	1047
	Std. Dev. [kcal]	N/A	129	658
	Recipe count	N/A	17k	10k
	Sample count	N/A	**44k**	**24k**
pic2kcal	Mean [kcal]	425	179	1791
	Std. Dev. [kcal]	207	73	1007
	Recipe count	42k	70k	63k
	Sample count	**179k**	**308k**	**267k**

Javier Marin, Aritro Biswas, Ferda Ofli, Nicholas Hynes, Amaia Salvador, Yusuf Aytar, Ingmar Weber, and Antonio Torralba. Recipe1M+: A Dataset for Learning Cross-Modal Embeddings for Cooking Recipes and Food Images. pages 1–1.

Multi-task Prediction

- End-to-end nutrition and ingredient estimation from food images

- Architecture:
  - backbone: DenseNet, ResNet, pre-trained on ImageNet
  - last layer adapted for
    - regression outputs on kcal and macronutrients
    - binary outputs on top 100 ingredients

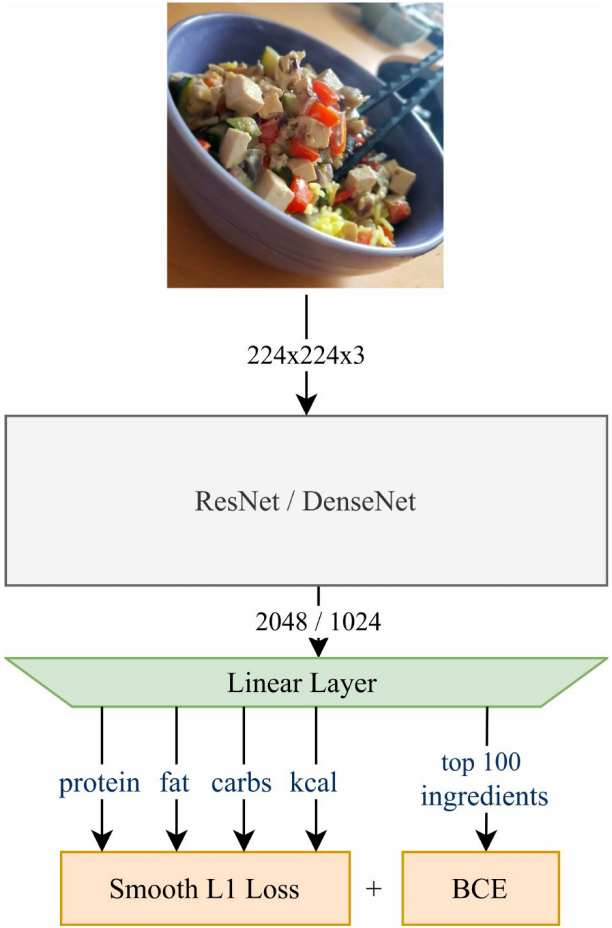

Multi-Task Learning for Calorie Prediction on a Novel Large-Scale Recipe Dataset Enriched with Nutritional Information, Ruede et al.

Multi-task Prediction

- Training:

$$\text{multi-task loss} = \text{L1}_{\text{kcal}} + \sum_{m \in \{\text{fat,prot,carb}\}} \text{L1}_m + \gamma \cdot \text{BCE}$$

- Evaluation:
  - Comparison against mean baseline
  - relative error for calories
  - absolute error for calories and macronutrients

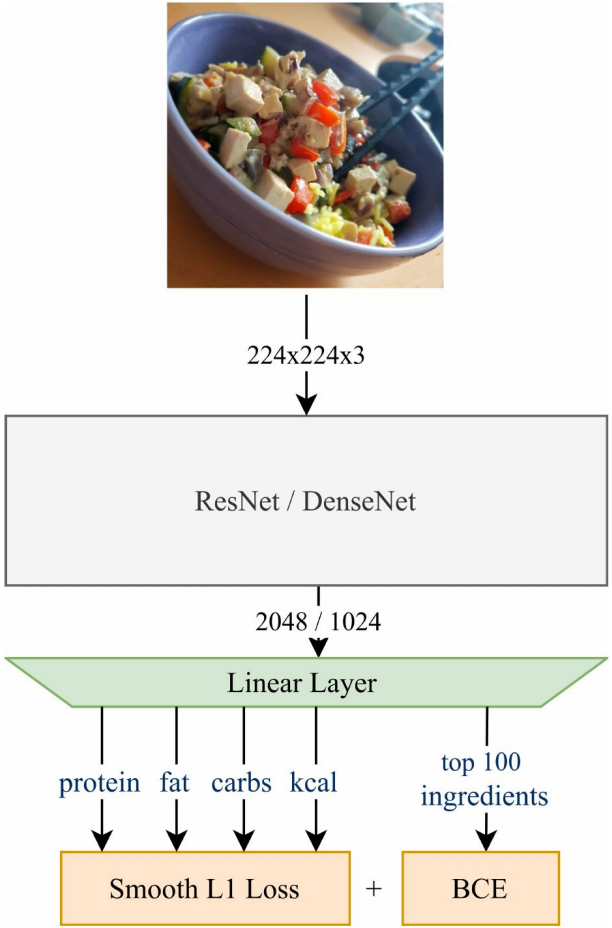

Multi-Task Learning for Calorie Prediction on a Novel Large-Scale Recipe Dataset Enriched with Nutritional Information, Ruede et al.

Results

- Relative and absolute error depending on the amount of food

	amount	kcal (rel)	kcal	protein	fat	carbs
Mean BL	portion	0.736	170	11.2	11.4	22.2
Ours		**0.623**	154	9.21	10.7	19.1
Mean BL	recipe	1.23	858	41.9	54.4	125
Ours		**0.823**	711	34.8	46.9	94.4
Mean BL	100g	0.464	60.5	3.10	4.49	10.5
Ours		**0.326**	46.9	2.51	3.88	6.97

- Relative and absolute error by prediction task

	kcal (rel)	kcal	protein	fat	carbs
Random Baseline	0.595	83.3	4.36	6.32	15.0
Mean Baseline	0.464	60.5	3.10	4.49	10.5
Kcal-only	0.362	50.3	N/A	N/A	N/A
∔ macros	0.345	49.0	2.67	4.06	7.70
∔∔ top-100 ingredients	**0.326**	46.9	2.51	3.88	6.97

Multi-Task Learning for Calorie Prediction on a Novel Large-Scale Recipe Dataset Enriched with Nutritional Information, Ruede et al.

Examples

Nutrition Facts (per 100 g)		
	Pred	True
Calories	183 kcal	198 kcal
Fat	9 g	9 g
Carb	17 g	24 g
Protein	7 g	4 g
Ingredients (pred): Flour, Butter, Milk		
Ingredients (true): Eggs, Flour, Vanilla Sugar		

Nutrition Facts (per 100 g)		
	Pred	True
Calories	229 kcal	239 kcal
Fat	3 g	2 g
Carb	44 g	46 g
Protein	7 g	7 g
Ingredients (pred): Flour		
Ingredients (true): Oil, Flour		

Nutrition Facts (per 100 g)		
	Pred	True
Calories	99 kcal	59 kcal
Fat	8 g	4 g
Carb	7 g	5 g
Protein	3 g	1 g
Ingredients (pred):		
Ingredients (true): Garlic		

Conclusion and Future Work

- Large dataset of 308k images with structured metadata
  - Generation code public https://github.com/phiresky/pic2kcal
- Ingredients, macronutrients, and calories are intertwined
  - Predicting them together improves performance
- Ingredient matching imperfect
- Extension to other tasks possible
  - Predicting the dietary style (e.g. vegan), “healthiness”

Multi-Task Learning for Calorie Prediction on a Novel Large-Scale Recipe Dataset Enriched with Nutritional Information, Ruede et al.