Controlled Recipe Generation: Adapting Food Recipes to Meet Dietary Restriction

Kemalcan Bora

Universitat Pompeu Fabra, Barcelona, Spain

Abstract

To address the growing demand for personalized nutrition in the face of rising rates of obesity, diabetes and heart disease, this paper proposes a controlled text generation framework designed to produce food recipes that meet specific dietary restrictions and nutritional goals. Our approach uses a comprehensive database, the NutriCuisine Index, containing 23,932 recipes with detailed dietary classifications, and transformer based models for dietary classification and nutrient estimation. Experimental results demonstrate robust performance, with a BERT based model achieving a macro F1 score of 0.94 for multi-label diet classification and a T5-3B model, equipped with a custom regression layer, achieving a R^2 of 0.913 for predicting nutrient content (carbohydrate, protein, fat and water). An optimization module adjusts ingredient quantities to meet user defined nutritional goals, while a sequence-to-sequence model generates cooking instructions. This study presents a framework for generating recipes that meet individual dietary and nutritional requirements.

Keywords

Language Models, Multi-Label Classification, Recipe Database, Controlled Text Generation

1. Introduction

The growing incidence of diet related illnesses, such as obesity, diabetes and heart disease, has increased the demand for personalized nutrition advice [1, 2, 3]. Non-communicable diseases claim millions of lives each year with heart disease responsible for 17.9 million deaths, cancer for 9.3 million, respiratory conditions for 4.1 million and diabetes for 2.0 million according to the WHO[1]. Diets that are customized to meet specific health needs or personal preferences like gluten-free diets for people with celiac disease, vegan diets for those with ethical or health motivations, nut-free diets for allergy sufferers and low-sugar diets for individuals managing blood sugar tend to be followed more consistently than standard diet plans that aren't tailored to anyone in particular. This is because personalized diets better match an individual's unique situation, making them easier and more appealing to stick with [4]. To support such adapted eating plans, precise assessment of carbohydrates, protein, fat and water is essential for managing energy, tissue repair, cardiovascular health and hydration [5, 6, 7, 8].

Public food datasets such as Recipe1M+ [9], USDA FoodData Central [10] and RecipeNLG [11] have driven advances in recipe classification, information extraction and generation. However, most recipe collections lack comprehensive diet labels and many recent generators ignore dietary constraints, risking unhealthy suggestions [12, 13, 14]. To address this, we propose a controlled text generation framework that produces ingredient lists and cooking steps aligned with user specified diet types and nutritional goals [15, 16]. We outline a five stage process for the proposed system:

- **Database construction**: Ingredients, their nutrient profiles and supported diet types are collected.
- **Diet-type classifier**: Pretrained language models are fine-tuned to assign one or more diet labels by ingredients.
- **Nutrient estimator**: A regression layer is developed to estimate the quantities of carbohydrates, protein, fat and water in grams based on ingredient list inputs.

Doctoral Symposium on Natural Language Processing, 25 September 2025, Zaragoza, Spain.
kemalcan.bora01@estudiant.upf.edu (K. Bora)
0009-0007-3607-2136 (K. Bora)

[1] https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases

- **Quantity controller**: Sequential Least Squares Programming is used to adjust ingredient amounts so predicted nutrients match user targets while respecting availability and diet rules.
- **Instruction generator and validator**: Step by step cooking directions are generated with a sequence-to-sequence model then the final ingredient list is reclassified to ensure it meets the chosen diet type with iteration if necessary.

Hypothesis

Driven by the research objectives, the following hypotheses are proposed to guide this research on developing a Controlled Text Generation (CTG) system:

- **H1**: A fine-tuned transformer based model (e.g., T5, BERT, or ALBERT) can perform accurate multi-label classification of ingredient lists into specific dietary categories (e.g., gluten-free, vegan, nut-free, low-sugar) enabling reliable identification of diet compliant recipes.
- **H2**: A regression-based nutrient estimation model, built on a fine-tuned language model (e.g., T5, BERT) can effectively predict the grams of carbohydrates, protein, fat and water from ingredient lists based on ground truth nutritional data.
- **H3**: A quantity optimization algorithm can successfully adjust ingredient amounts to meet user specified nutritional targets (e.g., grams of carbohydrates, protein, fat and water) while adhering to dietary restrictions.
- **H4**: A sequence-to-sequence model, combined with a validation mechanism, can produce coherent and diet compliant cooking instructions of high quality, ensuring usable recipe outputs.

These hypotheses will guide the research in creating a CTG system that addresses the growing need for health conscious culinary solutions by ensuring accuracy, nutritional alignment and practical applicability.

2. Related Works

The field of recipe generation has seen steady progress toward creating systems that follow regional styles and dietary needs. Kazama et al. were the first to identify mixtures of regional styles in recipes and to use an LSTM model to generate cooking instructions for each style [17]. Pan et al. measured how closely recipes match regional patterns, for example using the Mediterranean diet score [18].

Blackstone et al. compared US, Mediterranean and vegetarian diets to show how food choices affect health and the environment [19]. Mohammadi et al. added simple language features to neural models and saw big gains in classifying recipe difficulty [20].

Early work by Dale and Reiter built the EPICURE system, which used an ontology and grammar to describe cooking steps in recipes rather than invent new dishes [21]. The Transformer model then became the backbone of recipe generation. Petroni et al. showed that models like BERT store factual knowledge without extra training [22]. Today, models such as T5, BART and GPT drive most data to text tasks. Yin and Wan tested different sequence-to-sequence models on benchmarks like E2E, WikiBio, WebNLG and ToTTo and found that fine-tuned transformers score higher on BLEU [23].

In controlled text generation, researchers aim to satisfy specific constraints while maintaining output diversity. Zhang et al. utilized adversarial training to enhance the variety of generated text by matching high-dimensional latent feature distributions of real and synthetic sentences, thereby addressing the issue of limited output variety [24]. Ke et al. proposed Adversarial Reward Augmented Maximum Likelihood (ARAML) for more stable training [25]. Zhou et al. created Controlled Text Generation (CTG), which applies natural language rules with in context learning [26]. Pascual et al. built a plug and play recipe generator with content planning to meet diet needs, showing how a plug and play approach can handle restrictions like low salt or vegan diets [15]. Lee et al. used contrastive learning to cut down on biased outputs in recipe generation [27] and Jie et al. showed that soft prompt tuning on T5 lets models control attributes like sentiment or style [28].

For diet label classification, Adhikari et al. used BERT with knowledge distillation to handle long recipe texts with multiple labels [29]. Other work applied support vector machines and enhancements to LIBLINEAR for fast and accurate sorting of recipes by dietary labels [30, 31]. Pranesh and Shekhar explored small models that run well on limited hardware [32].

All these studies demonstrate that classification methods can produce recipes that follow regional styles and meet dietary needs.

3. Proposed Methodology

This paper proposes a methodology centered on framework for Controlled Text Generation (CTG), specifically applied to recipe generation. The approach comprises several key stages designed to produce health conscious recipes adhering to specific dietary requirements and nutritional targets.

Initially, we construct the NutriCuisine Index, a specialized database building upon foundational datasets like RecipeNLG [11]. Our index introduces detailed dietary classifications. Each entry contains the recipe title, servings, ingredient list with quantities, cooking instructions and assigned diet labels (e.g., Vegan, Gluten-Free). This extension of RecipeNLG provides enriched data crucial for training the following models.

Next, we develop a multi-label diet type classification model by fine-tuning transformer based models (e.g., T5, ALBERT, BERT) on the NutriCuisine Index. These models learn to predict relevant diet labels (e.g., High-Protein, Nut-Free) from an ingredient list. To handle potential class imbalance, model evaluation prioritizes the F1-score (micro, macro, weighted). Standard fine-tuning techniques like learning rate optimization and early stopping are employed. The resulting classifier verifies compliance of generated recipes with user specifications.

Subsequently, we address nutrient estimation. A custom regression layer is integrated into a fine-tuned T5 model to predict nutritional content (grams of carbohydrates, protein, fat, water) from ingredient lists. Trained on the NutriCuisine Index using ground truth data from USDA FoodData Central [10], the regression head outputs continuous nutrient values. Performance is evaluated using Mean Squared Error (MSE) and compared against a naive baseline to demonstrate effectiveness.

A key control mechanism is quantity optimization. We implement a Sequential Least Squares Programming (SLSQP) [33] optimizer to adjust ingredient quantities, directly controlling the recipe's nutritional profile to meet user targets for macronutrients. It utilizes predictions from the T5 regression model and modifies quantities, potentially removing ingredients if necessary to satisfy nutritional constraints. This process incorporates realistic quantity limits based on servings and typical ingredient availability, ensuring the input for text generation strictly adheres to numerical and compositional controls.

Finally, the methodology includes recipe generation and validation. Using the optimized ingredient list from the SLSQP step, a fine-tuned T5 model generates step-by-step cooking instructions via a prompt based sequence-to-sequence approach. The prompt includes the detailed ingredient list and target diet type(s). A critical validation step follows: the generated recipe's ingredients are processed by the diet type classifier. If predicted labels match the user's specification, the recipe is considered compliant; otherwise, the process can iterate. Recipe quality is assessed using metrics (BLEURT [34], ROUGE [35]) and potentially supplemented by expert evaluation [36, 37, 14].

This comprehensive methodology leverages transformer models and optimization techniques for the systematic generation of health conscious recipes meeting specific dietary and nutritional constraints.

4. Experiments

This section details the experimental procedures and results related to the construction of the NutriCuisine Index and the development of the nutrient estimation model.

4.1. NutriCuisine Index Construction and Characteristics

The first experimental step involved the creation of the NutriCuisine Index. We compiled a dataset comprising *23,932* recipes sourced from publicly available websites, including BBC Good Food, Heart UK and Delish. A thorough review confirmed compliance with special data regulations and the General Data Protection Regulation (GDPR), permitting the use of this data for research purposes. These sources provided essential recipe information, including ingredients, preparation steps and initial nutritional and dietary details.

A key contribution of the NutriCuisine Index is its focus on dietary classifications, addressing a gap present in existing datasets like RecipeNLG which often lack explicit diet type information. Our database includes both multi-label and single-label classifications. The diet labels were established through a two stage process: Initial labels were collected during web scraping, followed by expert validation performed by two commissioned dietitians. These experts reviewed each recipe, verifying existing labels and adding new ones based on professional assessment of ingredients and nutritional content. The final database schema for NutriCuisine encompasses fields such as Title, Serve, Link, Ingredients (with quantities), Directions, Nutrition and the validated Diets list, providing a comprehensive overview for each recipe (detailed in Table 1).

Table 1
NutriCuisine Database Schema and Dietary Type Counts

Database Schema

Field	Description
Title	Recipe name
Serve	Number of servings
Link	URL to original recipe source
Ingredients	List of ingredients with quantities
Directions	Step by step cooking instructions
Nutrition	Nutritional content per serving
Diets	Suitable diet types for the recipe

Dietary Type Counts

Dietary Type	Count	Dietary Type	Count
30-Minute-Meals	129	**Low-Carb**	**1469**
Appetizers	115	**Low-Sugar**	**1556**
Dairy-Free	215	Lunch	58
Dinner	158	**Nut-Free**	**2372**
Easily-Doubled	597	One-Pot-Meals	86
Easily-Halved	447	Pressure-Cooker	45
Freezable	4096	Salads	61
Gluten-Free	**11353**	Slow-Cooker	63
Healthy	**9416**	Soup	60
High-Protein	**864**	**Vegan**	**20438**
Kid-Friendly	99	Whole-30	74

4.2. Multi-Label Diet Classification

This experiment focused on classifying recipes from the NutriCuisine Index into seven key dietary categories: **Gluten-Free**, **Healthy**, **High-Protein**, **Low-Carb**, **Low-Sugar**, **Nut-Free** and **Vegan** using transformer based models trained directly on ingredient text.

Data Preparation: Ingredient lists sourced from the NutriCuisine Index were preprocessed to ensure consistency and reduce noise. This involved numeric standardization (e.g., converting fractions to decimals), text cleaning (including punctuation removal and Unicode normalization) and unit standardization. The corresponding diet labels for the seven target categories were binarized using MultiLabelBinarizer for multi-label classification.

Model Architecture and Training: Four transformer based models were adapted for this multi-label classification task these are *BERT-Base-Uncased, RoBERTa-Base, ALBERT-Base-V2* and *DistilBERT-Base-Uncased*. The output layer of each model was configured to predict probabilities for the seven dietary categories. The dataset was partitioned into 70% for training and 30% for testing. Models were trained for up to 5 epochs using a batch size of 8. AdamW optimizer [38] was employed with a learning rate of 1e-5 and binary cross-entropy was used as the loss function. EarlyStopping with a patience of 3 epochs and a minimum delta of 0.02 was implemented to prevent overfitting.

Evaluation: Model performance was assessed using standard multi-label classification metrics (precision, recall and F1-score) To account for potential class imbalance among the dietary categories, results were reported using micro, macro and weighted averaging across the seven classes.

Results and Analysis: The performance of the models trained is detailed in Table 2.

Table 2
Performance of Models: F1-Scores

Class	BERT			RoBERTa			ALBERT			DistilBERT		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Gluten-Free	0.98	1.00	**0.99**	0.95	0.99	0.97	0.88	0.99	0.94	0.99	0.98	0.98
Healthy	0.98	0.97	**0.98**	0.93	0.98	0.96	0.92	0.97	0.94	0.98	0.98	0.98
High-Protein	0.95	0.97	0.96	0.90	0.91	0.90	0.92	0.76	0.84	0.95	0.97	**0.96**
Low-Carb	0.99	1.00	**0.99**	0.98	0.99	0.99	0.98	0.99	0.98	0.99	0.99	**0.99**
Low-Sugar	1.00	0.56	**0.72**	1.00	0.26	0.41	1.00	0.23	0.38	1.00	0.28	0.44
Nut-Free	0.98	0.97	**0.97**	0.94	0.97	0.95	0.93	0.86	0.89	0.94	0.98	0.96
Vegan	0.99	0.99	**0.99**	0.96	1.00	0.98	0.99	0.94	0.96	0.98	1.00	**0.99**
Macro Avg.	0.98	0.92	**0.94**	0.95	0.87	0.88	0.95	0.82	0.85	0.97	0.88	0.90

Overall, the models demonstrated strong classification capabilities on ingredients. BERT-Base-Uncased achieved the highest macro-averaged F1-score at 0.94, indicating reliable performance across the different diet types. DistilBERT-Base also performed well with a macro F1 of 0.90, followed by RoBERTa-Base (0.88) and ALBERT-Base-V2 (0.85).

Examining individual class performance reveals high F1-scores (often 0.98-0.99) for categories like Gluten-Free, Low-Carb and Vegan across most models, suggesting these diets have distinct ingredient patterns that are well captured. Healthy and Nut-Free also generally showed strong results (F1 typically >0.94). High-Protein classification was solid, though slightly less consistent across models compared to the top performers. The Low-Sugar category proved most challenging, particularly in terms of recall, resulting in lower F1-scores compared to other categories (e.g., BERT achieved 0.72 F1, while others were lower). These results highlight the effectiveness of transformer models for dietary classification directly from ingredient lists, while also identifying specific categories that remain more difficult to predict accurately based solely on text.

4.3. Nutrient Estimation Model: Setup and Results

For the nutrient estimation task, we developed and evaluated a T5Regressor model. This model adapts the encoder component of pretrained T5 models (specifically testing T5-small, T5-base, T5-large and T5-3B variants) for regression. The architecture uses the T5 encoder to generate contextual embeddings from input food names. The encoder's last hidden state is mean-pooled across the sequence length (weighted by the attention mask) to produce a fixed size representation. This representation is then fed into a sequential head consisting of a dropout layer (rate=0.2) and a linear layer, which outputs four continuous values corresponding to the target nutrients: carbohydrates, protein, fat and water (in grams).

The training and evaluation were performed using the USDA FoodData Central dataset, containing 7,793 food entries with descriptive names and corresponding nutrient values. Preprocessing involved tokenizing the food names using the appropriate T5 tokenizer for each model variant. Analysis showed an average token length of 12 tokens (max 32), leading us to set a maximum sequence length of 150 tokens to avoid truncation while managing computational load; shorter sequences were padded. The target nutrient values were standardized using StandardScaler to achieve zero mean and unit variance, aiding training stability.

The dataset was split into 80% for training (~6,234 samples) and 20% for testing (~1,559 samples) using a fixed random seed (42) for reproducibility. Training utilized dataloaders with a batch size of 32 and shuffling enabled for the training set. Model performance was evaluated using Mean Squared Error (MSE) and R^2 score, comparing against a naive baseline that predicts the mean nutrient value for all foods in the test set.

Model performance was evaluated using Mean Squared Error (MSE), which also served as the loss function during training, and R^2 score. Additionally, we report Mean Absolute Error (MAE) and Median Absolute Error (MDAE) to assess prediction quality from complementary perspectives. MAE quantifies the average prediction error in absolute terms, while MDAE is more robust to outliers, capturing the median absolute error. A naive baseline was also included, which simply predicts the mean nutrient value of the training set for all test samples.

Results and Analysis: The fine-tuned T5 models demonstrated effective learning for nutrient prediction. As shown in Table 3, all T5 variants significantly outperformed the naive baseline. Performance scaled directly with model size: T5-small achieved an R^2 of 0.648 and MSE of 139.84, while the largest model, T5-3B, yielded the best results with an R^2 of 0.913 and an MSE of 40.87, indicating it could explain approximately 91.3% of the variance in the true nutrient values. T5-large also performed strongly (R^2 0.894, MSE 47.63), offering a compelling balance between performance and model size, while T5-base performed intermediately. The "Loss" column in Table 3 reflects the final validation loss (MSE) during training.

Table 3
Overall performance comparison of fine-tuned T5 models, sorted by R^2

Model	R^2	Loss	MAE	MDAE	MSE
T5-3B	0.913	0.0932	3.16	1.60	40.87
T5-large	0.894	0.1130	3.55	1.85	47.63
T5-base	0.838	0.1740	4.43	2.41	65.99
T5-small	0.648	0.3835	6.90	4.08	139.84

Analysis of predictions for individual nutrients (detailed in Table 4) revealed that protein consistently had the lowest MSE across all models, suggesting it was the easiest nutrient to predict from food names. Conversely, water exhibited the highest MSE, indicating greater prediction difficulty. Carbohydrates and fat fell in between, with carbohydrates generally showing slightly higher MSE than fat. Importantly, even the smallest T5 model substantially improved upon the baseline for all nutrients (e.g., T5-3B reduced carbohydrate MSE from the baseline's 649.65 to 52.30), confirming that the models learned meaningful patterns from the food names related to nutritional content.

Table 4
Detailed MSE comparison for fine-tuned T5 models and the naive baseline across nutritional components. Lower values are better.

Model	MSE			
	Carbohydrate	Protein	Fat	Water
T5-3B	52.30	11.72	26.66	72.81
T5-large	57.91	15.71	32.84	84.08
T5-base	84.53	30.07	46.18	103.16
T5-small	152.34	52.99	147.70	206.34
Naive Baseline	649.65	125.24	283.90	897.33

5. Points for Further Discussion

While our controlled recipe generation system yields promising outcomes, still have a long way to go to develop and complete the framework. First, we plan to expand the NutriCuisine Index to include a broader diversity of cuisines and diet types, such as regional emerging dietary trends, to improve the system's applicability and generalization across varied culinary contexts. Second, integrating user feedback mechanisms such as ratings for taste, feasibility, or ingredient preferences could refine the personalization of generated recipes, making them more responsive to individual needs. Apart from this, we are still working on the methodology we propose, which is to control and optimize the nutritional values according to the desired amounts and to transform the obtained results into a recipe through a language model.

Beyond these planned improvements, several open questions invite further investigation:

- **Cultural and Regional Adaptation**: How can the system effectively incorporate cultural and regional culinary variations while ensuring compliance with dietary restrictions?
- **Transparency and Explainability**: What approaches can be developed to make the recipe generation process more interpretable, such as explaining ingredient selections or quantity adjustments to foster user trust and engagement?
- **Real World Validation**: How do the generated recipes perform in practical settings? Comprehensive evaluations with diverse user groups are needed to assess taste, preparation feasibility and nutritional adequacy compared to human crafted recipes.

Addressing these challenges and questions will be critical to advancing controlled recipe generation, ultimately enabling the delivery of highly personalized, health focused culinary solutions that meet both practical and nutritional demands.

Data and Software Availability

The datasets and source code supporting the findings and methodology presented in this study are publicly available to ensure reproducibility and encourage further research. The specific resources include:

- **Multi-Label Diet Classification Code:** The implementation of our diet classification model can be found at: https://github.com/NutriCuisine/NERonLLM
- **Nutrient Estimation Model Code:** The source code for the nutrient estimation component is available at: https://github.com/NutriCuisine/NutrientsFinder
- **NutriCuisine Index:** The dataset developed for this work is hosted at: https://github.com/NutriCuisine/database
- **FoodData Central Dataset:** We utilized the publicly available USDA FoodData Central SR Legacy dataset (April 2018 release, JSON format) for foundational nutrient information, accessible via: https://fdc.nal.usda.gov/fdc-datasets/FoodData_Central_sr_legacy_food_json_2018-04.zip

Declaration on Generative AI

The author have not employed any Generative AI tools.

References

[1] W. P. T. James, T. Gill, Obesity–introduction: history and the scale of the problem worldwide, Clinical Obesity in Adults and Children (2022) 1–16.

[2] K. Voigt, S. G. Nicholls, G. Williams, Childhood obesity: ethical and policy issues, Oxford University Press, USA, 2014.
[3] E. A. Reece, G. Leguizamón, A. Wiznitzer, Gestational diabetes: the need for a common ground, The Lancet 373 (2009) 1789–1797.
[4] T. Cruwys, R. Norwood, V. S. Chachay, E. Ntontis, J. Sheffield, "an important part of who i am": The predictors of dietary adherence among weight-loss, vegetarian, vegan, paleo, and gluten-free dietary groups, Nutrients 12 (2020) 970.
[5] D. J. Jenkins, C. W. Kendall, L. S. Augustin, S. Franceschi, M. Hamidi, A. Marchie, A. L. Jenkins, M. Axelsen, Glycemic index: overview of implications in health and disease, The American journal of clinical nutrition 76 (2002) 266S–273S.
[6] J. Bauer, G. Biolo, T. Cederholm, M. Cesari, A. J. Cruz-Jentoft, J. E. Morley, S. Phillips, C. Sieber, P. Stehle, D. Teta, et al., Evidence-based recommendations for optimal dietary protein intake in older people: a position paper from the prot-age study group, Journal of the american Medical Directors association 14 (2013) 542–559.
[7] R. J. De Souza, A. Mente, A. Maroleanu, A. I. Cozma, V. Ha, T. Kishibe, E. Uleryk, P. Budylowski, H. Schünemann, J. Beyene, et al., Intake of saturated and trans unsaturated fatty acids and risk of all cause mortality, cardiovascular disease, and type 2 diabetes: systematic review and meta-analysis of observational studies, Bmj 351 (2015).
[8] B. M. Popkin, K. E. D'Anci, I. H. Rosenberg, Water, hydration, and health, Nutrition reviews 68 (2010) 439–458.
[9] J. Marin, A. Biswas, F. Ofli, N. Hynes, A. Salvador, Y. Aytar, I. Weber, A. Torralba, Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images, IEEE Trans. Pattern Anal. Mach. Intell. (2019).
[10] J. K. Ahuja, A. J. Moshfegh, J. M. Holden, E. Harris, Usda food and nutrient databases provide the infrastructure for food and nutrition research, policy, and practice, The Journal of nutrition 143 (2013) 241S–249S.
[11] M. Bień, M. Gilski, M. Maciejewska, W. Taisner, D. Wisniewski, A. Lawrynowicz, Recipenlg: A cooking recipes dataset for semi-structured text generation, in: Proceedings of the 13th International Conference on Natural Language Generation, 2020, pp. 22–28.
[12] Y. Pan, Q. Xu, Y. Li, Food recipe alternation and generation with natural language processing techniques, in: 2020 IEEE 36th International Conference on Data Engineering Workshops (ICDEW), IEEE, 2020, pp. 94–97.
[13] M. Bień, M. Gilski, M. Maciejewska, W. Taisner, Cooking recipes generator utilizing a deep learning-based language model, Taisner, Cooking recipes generator utilizing a deep learning-based language model (2020).
[14] A. Reusch, A. Weber, M. Thiele, W. Lehner, Recipegm: A hierarchical recipe generation model, in: 2021 IEEE 37th International Conference on Data Engineering Workshops (ICDEW), IEEE, 2021, pp. 24–29.
[15] D. Pascual, B. Egressy, C. Meister, R. Cotterell, R. Wattenhofer, A plug-and-play method for controlled text generation, arXiv preprint arXiv:2109.09707 (2021).
[16] B. Guo, H. Wang, Y. Ding, W. Wu, S. Hao, Y. Sun, Z. Yu, Conditional text generation for harmonious human-machine interaction, ACM Transactions on Intelligent Systems and Technology (TIST) 12 (2021) 1–50.
[17] M. Kazama, M. Sugimoto, C. Hosokawa, K. Matsushima, L. R. Varshney, Y. Ishikawa, A neural network system for transformation of regional cuisine style, Frontiers in ICT 5 (2018) 14.
[18] S. Pan, L. Dai, X. Hou, H. Li, B. Sheng, Chefgan: Food image generation from recipes, in: Proceedings of the 28th ACM International Conference on Multimedia, 2020, pp. 4244–4252.
[19] N. T. Blackstone, N. H. El-Abbadi, M. S. McCabe, T. S. Griffin, M. E. Nelson, Linking sustainability to the healthy eating patterns of the dietary guidelines for americans: a modelling study, The Lancet Planetary Health 2 (2018) e344–e352.
[20] E. Mohammadi, N. Naji, L. Marceau, M. Queudot, E. Charton, L. Kosseim, M.-J. Meurs, Cooking up a neural-based model for recipe classification, in: Proceedings of the Twelfth Language Resources

and Evaluation Conference, 2020, pp. 5000–5009.
[21] R. Dale, Cooking up referring expressions, in: 27th Annual Meeting of the association for Computational Linguistics, 1989, pp. 68–75.
[22] F. Petroni, T. Rocktäschel, S. Riedel, P. Lewis, A. Bakhtin, Y. Wu, A. Miller, Language models as knowledge bases?, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 2463–2473.
[23] X. Yin, X. Wan, How do seq2seq models perform on end-to-end data-to-text generation?, in: Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2022, pp. 7701–7710.
[24] Y. Zhang, Z. Gan, K. Fan, Z. Chen, R. Henao, D. Shen, L. Carin, Adversarial feature matching for text generation, in: International conference on machine learning, PMLR, 2017, pp. 4006–4015.
[25] P. Ke, F. Huang, M. Huang, X. Zhu, Araml: A stable adversarial training framework for text generation, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019, pp. 4271–4281.
[26] W. Zhou, Y. E. Jiang, E. Wilcox, R. Cotterell, M. Sachan, Controlled text generation with natural language instructions, in: Proceedings of the 40th International Conference on Machine Learning, 2023, pp. 42602–42613.
[27] S. Lee, D. B. Lee, S. J. Hwang, Contrastive learning with adversarial perturbations for conditional text generation, arXiv preprint arXiv:2012.07280 (2020).
[28] R. Jie, X. Meng, L. Shang, X. Jiang, Q. Liu, Prompt-based length controlled generation with reinforcement learning, arXiv preprint arXiv:2308.12030 (2023).
[29] A. Adhikari, A. Ram, R. Tang, W. L. Hamilton, J. Lin, Exploring the limits of simple learners in knowledge distillation for document classification with docbert, in: Proceedings of the 5th Workshop on Representation Learning for NLP, 2020, pp. 72–77.
[30] S. Ghosh, A. Dasgupta, A. Swetapadma, A study on support vector machine based linear and non-linear pattern classification, in: 2019 International Conference on Intelligent Sustainable Systems (ICISS), IEEE, 2019, pp. 24–28.
[31] G. Hinselmann, L. Rosenbaum, A. Jahn, N. Fechner, C. Ostermann, A. Zell, Large-scale learning of structure- activity relationships using a linear support vector machine and problem-specific metrics, Journal of chemical information and modeling 51 (2011) 203–213.
[32] R. Pranesh, A. Shekhar, Analysis of resource-efficient predictive models for natural language processing, in: Proceedings of SustaiNLP: Workshop on Simple and Efficient Natural Language Processing, 2020, pp. 136–140.
[33] W. Murray, F. J. Prieto, A sequential quadratic programming algorithm using an incomplete solution of the subproblem, SIAM Journal on Optimization 5 (1995) 590–640.
[34] T. Sellam, X. Du, A. Das, A. Parikh, Bleurt: Learning robust metrics for text generation, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, pp. 7881–7892.
[35] C.-Y. Lin, Rouge: A package for automatic evaluation of summaries, in: Text summarization branches out, 2004, pp. 74–81.
[36] W. Antô, J. R. Bezerra, L. F. W. Góes, F. M. F. Ferreira, et al., Creative culinary recipe generation based on statistical language models, IEEE Access 8 (2020) 146263–146283.
[37] D. Khashabi, G. Stanovsky, J. Bragg, N. Lourie, J. Kasai, Y. Choi, N. A. Smith, D. S. Weld, Genie: Toward reproducible and standardized human evaluation for text generation, arXiv preprint arXiv:2101.06561 (2021).
[38] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, in: International Conference on Learning Representations, 2019. URL: https://openreview.net/forum?id=Bkg6RiCqY7.