Detecting Changing Culinary Trends Through Historical Recipes

Gauri Bhagwat, Marieke van Erp, Teresa Paccosi, and Rik Hoekstra
DHLab, KNAW Humanities Cluster
Oudezijds Achterburgwal 185
1012 DK Amsterdam, the Netherlands
{gauri.bhagwat,marieke.van.erp,teresa.paccosi}@dh.huc.knaw.nl
rik.hoekstra@di.huc.knaw.nl

Abstract

Culinary trends evolve in response to social, economic, and cultural influences, reflecting broader historical transformations. We present an exploration into Dutch culinary trends from 1910 to 1995 by analysing recipes from housekeeping school cookbooks and newspaper recipe collections. Using computational techniques, we extract and examine ingredient frequency, recipe complexity, and shifts in recipe categories to identify trends in Dutch cuisine from a quantitative point of view. Additionally, we experimented with Large Language Models (LLMs) to structure and extract recipes' features, demonstrating their potential for historical recipe parsing.

1 Introduction

Globalisation and advanced production methods make our current food environment one of seemingly unlimited choice. However, what we eat has historically been dictated by availability of foodstuffs and socio-cultural traditions (Sonnenfeld et al., 1999). Historical food traditions were not static, for example colonisation has provided access to new, faraway flavours that found their way into local dishes (Protschky, 2008; Czarra, 2009). Additionally, societal events could influence the availability of certain goods, requiring cooks to adapt their recipes (Claflin, 2020). In this paper, we compare Dutch recipe sets from different points in time to detect changes in culinary trends. We do so by comparing different editions of Dutch housekeeping school manuals as well as recipes extracted from newspapers. While we do not know exactly which recipes were cooked most, these resources were widely read and used. We hypothesise that a possible indicator of a recipe or ingredient's declining popularity could be its absence (or decrease) from more recent recipe collections. Our contributions are: 1) a structured dataset of ingredients by year and recipe for Dutch historical housekeeping cookbooks, and 2) a diachronic analysis of Dutch culinary trends for the period 1910-1995. The remainder of this paper is organised as follows. In Section 2, we discuss related work on historical recipe extraction and analysis. In Section 3, we describe the datasets and the data processing, followed by our analyses (Section 4) and future directions (Section 5). Our code and data can be found at: https://github.com/trifecta-project/Dutch-historical-recipe-trends

2 Related Work

There is a fair body of research on analysing contemporary recipes for a variety of tasks such as personalised food recommendations (e.g. Teng et al., 2012; Haussmann et al., 2019; Jain and Singhal, 2022) or ingredient substitutions (e.g. Elsweiler et al., 2017; Shirai et al., 2021). Many of these approaches rely on existing recipe datasets such as the Recipe1M+ dataset (Marın et al., 2021) or exploit semantic markup in recipe websites that allows for easy extraction of ingredients and other types of information (Chatterjee et al., 2016). Historians have a long tradition of analysing historical recipes using qualitative methods rather than automated, quantitative ones (Leong, 2019). Efforts such as the Feeding America - the Historic American Cookbook Project have aimed to close this gap (Berg and Jones, 2003). In this project, 76 cookbooks were manually annotated with information on ingredients, cooking methods and additional categorisations per recipe. Advances in language technology and the large-scale digitisation of collections have made automatic historical recipe analysis possible (van Erp et al., 2018). In this work, we combine and extend (van Erp et al., 2018) to automatically extract ingredients and other characteristics from historical recipes and show how their changing over time illustrates changing culinary tastes.

Abbr.	Year	Abbr.	Period
AHS_1	1910	HRW_1	1946-50
AHS_2	1912	HRW_2	1951-60
AHS_3	1925	HRW_3	1961-70
KHB_4	1932	HRW_4	1971-80
CHB_5	1938	HRW_5	1981-90
MHS_6	1939	HRW_6	1991-95
AHS_7	1940	-	-

Table 1: Abbreviations for the cookbooks with publication year (columns 1-2) and Historical Recipe Web separated per selected periods (columns 3-4).

3 Data and Preprocessing

In this section, we describe the datasets used and the preprocessing steps we employed to refine and organise them for analysis.

3.1 Corpus Description

Our corpus consists of two data collections: seven historical cookbooks and a set of recipes published in newspapers.

Cookbooks (1910-1940)

This dataset includes four different editions of the *Kookboek van de Amsterdamsche Huishoudschool* (AHS), *Praktische recepten van de Huishoudschool Mariakroon Culemborg* (MHS), *Kook-en huishoudboek voor het platteland* (KHB), and *Het coöperatieve kook-en huishoudboek* (CHB). These cookbooks were designed as instructional guides for the domestic education of young girls attending housekeeping schools. The 'huishoud' or housekeeping schools played an important role, with 25% of all girls in the Netherlands attending such institutions in the final stage of their school life at a time when most women did not have secondary education at all (Verwey-Jonker, 1955). The cookbooks were not limited to the schools as they were also widely purchased by the general public (Meijer, 2014), and covered a wide range of recipes, from traditional Dutch dishes to more intercultural influences, making them a valuable source for studying Dutch cooking practices. The AHS and MHS served as standardised instructional texts for domestic culinary practices in the *Huishoudschools* (EN: Housekeeping Schools) of urban areas. The KHB includes recipes from both urban and rural contexts, while the CHB highlights cooperative household management and resource-efficient meal planning.

Historical Recipe Web (HRW) (1946-1995)

This dataset includes automatically extracted multilingual recipes from Dutch newspapers (van Erp et al., 2018). They consist of 27,411 Dutch recipes from 4 newspapers spanning almost 6 decades. For this study, only the Dutch-language recipes were included, resulting in a dataset of 26,662 recipes. These recipes are standalone entries aimed at offering variety to readers, reflecting contemporary culinary trends.

The two datasets span nearly a century but they differ in coverage. In our comparisons, we consider temporal trends only within each dataset rather than across them. This means that while we track changes over time within cookbooks and newspapers separately, our comparisons focus on domain differences rather than direct chronological overlap. This approach allows us to highlight how recipe complexity and ingredient usage evolve within their respective contexts, reflecting the distinct functions of cookbooks and newspapers. All data is downloaded from Delpher,[1] the National Library of The Netherlands online repository of digitised printed material from the Netherlands. The text is in historical Dutch and presents challenges such as old-fashioned spelling, inconsistent formatting, and OCR errors. As the newspaper recipes are already structured, the remainder of this section focuses on preprocessing the cookbooks.

3.2 Data Cleaning

For data cleaning we concentrated on noise removal. The original text contained several information irrelevant to the present study, such as metadata, housekeeping instructions, and non-recipe text. This study focuses on recipe analysis, thus sections including data such as author names and book titles, were filtered out using regular expression, leaving only the recipes for processing.

3.3 Recipe Extraction

After data cleaning, we extracted individual recipes by segmenting the texts into recipe numbers and titles, and filtered the instructions using regular expressions. The extraction success rate for each book was calculated as the percentage of correctly extracted recipes out of the total, yielding an accuracy of 98.5%. For instance, common OCR errors such as missing periods at the end of the text ("."),

[1] https://delpher.nl Last visited: 27 February 2025

Extraction Quality	Count
Similar	32
Original more complete	6
Original more content	6
LLM more content	5
Original parsing error	1

Table 2: Comparison of LLM parsing result with original text. A legend is provided in the GitHub repository.

or its replacement with commas (",") interfered with the extraction process. To address the missing 1.5%, we manually added the excluded recipes, not previously detected due to this lack of uniform structural cues.

3.4 Ingredient Extraction

Due to the challenges posed by the absence of a Dutch recipe parser, OCR errors, and the non-perfect grammatical nature of the data, we use a dictionary-based matching approach to extract ingredients. We adapted and cleaned the ingredient lexicon from (van Erp et al., 2018) by combining a list of seasonal ingredients mentioned in cookbooks with the HRW lexicon and filtering irrelevant data for the present study, such as non-food terms and product names. The final lexicon includes 1,843 terms. Ingredients which are not included in the dictionary are not considered, limiting the list of extracted items to our lexicon. As shown in table 3, we obtained a total of 1,530 unique ingredients from both the 33,416 recipes of the cookbooks and the newspaper recipes.

3.5 LLM-based Extraction

We also experimented with Large Language Models (LLMs) to parse the recipes into a more structured format. We tested a 50 recipe sample from the 1925 edition of the AHS, using the llama3.2 LLM model.[2] We first parsed the recipes to a structured representation, consisting of a list of ingredients and a list of cooking instructions which we compared to the original. In general, the LLM parsing was very good (table 2), yielding a correct JSON representation of the recipes containing both a list of ingredients and a list of processing steps. However, as in the rule-based parsing, the LLM could not make sense of some of the content of the original recipes, in particular due to OCR errors. This holds especially true in case of numerical quantities and abbreviations for measures, such as litres. In the original text, this is commonly abbreviated as l., often mistaken for 1, or the other way round. For a more reliable result, an initial data cleaning step is needed which was out of the scope of this work, thus we decided to leave further investigation of this methodology to future work.

[2] https://huggingface.co/meta-llama/Llama-3.2-1B

4 Analysis and Discussion

To examine shifts in culinary trends, we analysed possible changes in ingredient usage over time. We focused on four key aspects: 1) recipe categories, 2) ingredient usage patterns, 3) recipe complexity, and 4) seasonal usage.

4.1 Recipe Categories

The analysis of the recipe categories was conducted by evaluating different meal categories, and the number of recipes in each category across different cookbooks. This analysis revealed notable culinary shifts in the different AHS' editions. The increase of recipe number in categories like *"verschillende voorgerechten of hartige voorproefjes"* (EN: various starters) and *"tusschengerechten en twaalfuurschotels"* (EN:entremets and main dishes) might point towards more elaborate meal compositions and a possible shift in meal structure. In the 1940s, not only meal structure but also categories saw a notable change with the introduction of a new one, *"rijsttafel"* (EN: rice table) with 22 recipes. This reflects the growing influence of South-East Asian Dutch colonies (now Indonesian) cuisine, likely a result of the popularity among Dutch people returning from the colonies. Between 1925 and 1940, another notable change is the decline in number of *"warme puddingen"* (EN: warm puddings) which decreased from 18 to 15, while *"koude nagerechten"* (EN: cold desserts) saw a significant increase, rising from 65 recipes in 1925 to 77 in 1940. The popularity of sweet categories such as *"gedroogde vruchten en compotes"* (EN: dried fruits and compotes) and *"zoete sausen"* (EN: sweet sauces) shows a consistent demand for sweetness. However, the decline of *"zure en zoete geleien"* (EN: sweet and sour jellies) in the 1940s, suggests a possible shifting attitude in sweet consumption. These changes, alongside the introduction of new categories like *"rauwkost"* (EN: raw vegetable dishes) and innovative preservation methods, may indicate a growing interest in health-related habits, while

Dataset	Year	Recipe Count	Unique Ingredients
Cookbooks	1910-1940	6,754	506
Historical Recipe Web	1946-1995	26,662	1,024
Total	1910-1995	33,416	1,530

Table 3: Recipe Count and Unique Ingredients for the Datasets

also reflecting colonial culinary influences, as well as possible war-induced rationing.

4.2 Ingredient Usage Pattern

We analysed changes in ingredient usage by looking at the 10 most frequent ingredients shared by both datasets and their oscillations across time. The heat map (Fig. 2) provides a visualisation of these ingredients, together with the 5 more frequent ones specific to each dataset. [3] The color scale in the legend ranges from light yellow (low frequency) at 0.0, indicating the absence of an ingredient, to deep red (high frequency) at 0.7, meaning that the ingredient appears in 70% of recipes. This gradient effectively illustrates variations in ingredient prevalence over time. Up until 1940, staple ingredients such as butter, eggs, sugar, and milk remained consistently popular, reflecting the didactic nature of cookbooks focused on essential ingredients for various recipes. Comparing the two datasets, changes in ingredients usage are more evident in the newspaper recipes. The higher use of *margarine* in just 1946-50 (HRW_1) possibly suggests its increase as a post-war alternative to heavily taxed ingredients. The heatmaps (Fig. 2) show a decline in *suiker* (EN: sugar) use over time, which may correspond to the similar trend we observed in the recipe categories. Meanwhile, the growing use of spices and herbs like *peper* (EN: pepper), and *peterselie* (EN: parsley) hints they transitioned from luxury items to everyday staples, potentially due to increased accessibility, evolving consumer preferences, or broader culinary influences.

4.3 Recipe Complexity

Recipe complexity can be assessed using various criteria, such as the number of preparation steps, required techniques, or ingredient diversity (Arendholz et al., 2013). Our two datasets (cookbooks and newspapers) are distinct in their purpose and context. The former served as instructional tools aimed at teaching traditional Dutch cuisine, while the latter is characterised by the novelty of its recipes, designed to capture reader's curiosity and inspire culinary experimentation. In this study, we casted ingredient frequency as a measure of complexity, with the assumption that a higher number of ingredients corresponds to a greater complexity of the recipe. While the cookbook recipes remained relatively consistent in the number of ingredients, the newspaper recipes exhibited a higher variance, with some recipes containing more than 50 ingredients. To investigate possible changes in the frequency of ingredients over time, we aggregated the total number of ingredients in each newspaper recipe by decade, as using decades would better capture changes over time compared to a yearly analysis. We calculated then the mean ingredient count for each decade to assess possible changes in ingredient usage (Fig. 1). To further explore variations in recipe complexity, we conducted a Pearson correlation to determine the statistical significance of the relationship between the rising frequency of ingredients in newspaper recipes and time. The results showed a positive correlation between the two variables, with a p-value < 0.005. We provide a more detailed visualisation of the distribution of ingredient counts across the two datasets in the boxplot (Fig. 3). The average number of ingredients in cookbooks remained stable with minimal fluctuations, reflecting a standard recipe pattern over time. From 1946-50, newspaper recipes showed

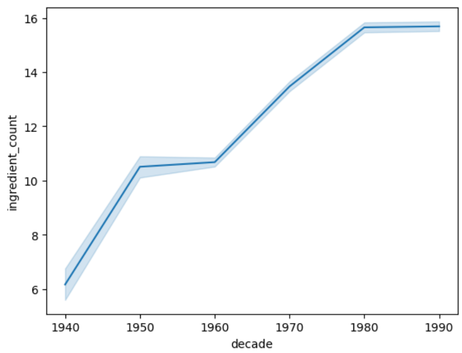

Figure 1: Mean count of ingredients over time and confidence interval in the newspaper dataset

[3]Ingredient EN translations are provided in the GitHub repository.

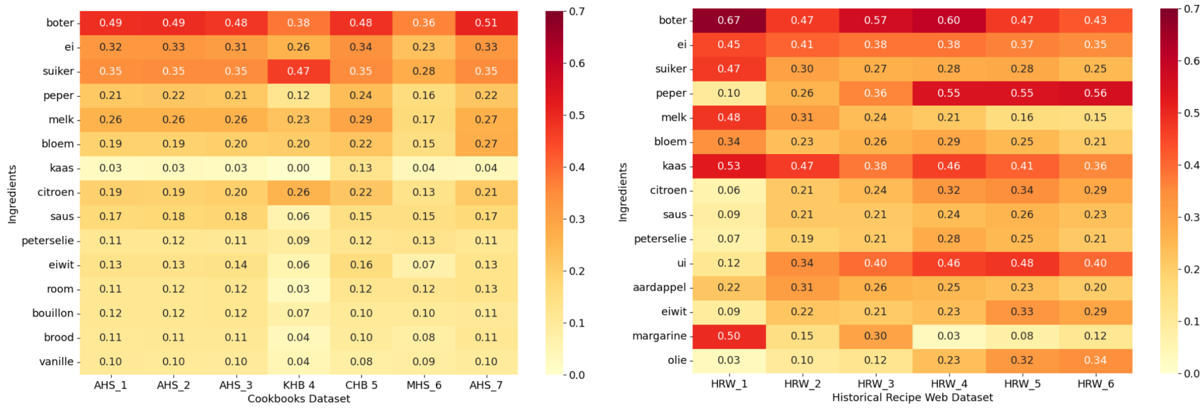

Figure 2: Heat map visualization of ingredient frequency for both datasets.

instead a rising trend as well as a growing number of outliers, reaching significantly higher complexity by the 1990s. This increase could be linked to the growing availability of non-traditional ingredients and rising wealth, particularly after the late 1950s. As economies grew and global trade expanded, supermarkets began offering a wider variety of products, including previously exotic ingredients. Additionally, increased international travel and immigration after the 1960s exposed people to foreign cuisines, making diverse ingredients more available. The increase in newspaper recipes may reflect an effort to educate home cooks on how to incorporate these new ingredients into their meals.

4.4 Seasonal Usage

The cookbooks provide a detailed list of ingredients for each month highlighting fruits, vegetables, meat, and fish, which are seasonally available. To compare seasonal ingredient use in newspaper recipes, we analysed ingredient frequency by season. For instance, since Dutch summer season runs from June to August, we identified summer ingredients from cookbooks and checked their presence in newspaper recipes for the same months. The results revealed that while the newspaper recipes were indeed more complex, they still promoted the use of seasonal ingredients. We found, for instance, that in summer months, more than 60 percent of newspaper recipes contained at least one ingredient from the traditional cookbook ingredient list. While this finding is not entirely unexpected, as the introduction of new ingredients did not lead to the disappearance of traditional ones, it is notable that some newspaper recipes featured more than 10 ingredients matching those found in the traditional cookbook lists for the same month. These results

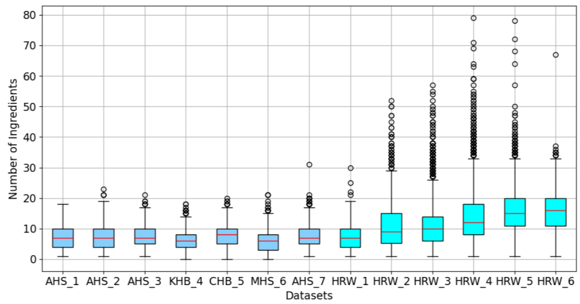

Figure 3: Distribution of ingredients counts in the cookbooks and newspaper datasets

suggest that although newspapers tended to publish novel recipes, they did not completely break away from the traditional culinary framework established by the cookbooks.

5 Conclusions and Future Work

This study provides a preliminary quantitative analysis of the evolution of culinary trends in Dutch recipes, highlighting how ingredient usage, recipe complexity, and culinary categories changed throughout the 20th century. Our findings indicate a trend toward increased recipe complexity in terms of larger number of ingredients, while still adhering to traditional culinary practices such as the use of seasonal ingredients. By releasing our dataset, we aim to support further research into culinary trends, which could be compared with other resources indicating ingredient availability due to economic factors or rationing constraints, as well as technological and climatological changes that may influence crop yields. Additionally, future work could refine our analyses by considering, for instance, variables beyond time as predictors for ingredient usage trends, as well as improving ingredient extraction methodologies to address OCR challenges, while expanding our dataset. As demonstrated by our preliminary experiments, the integration of LLMs could enhance the structural parsing of historical recipes, providing deeper insights into the evolution of culinary practices over time. In the future, we plan to address these outstanding points to further investigate culinary trends as a way of advance our understanding on the shaping of food and nutrition identity in Dutch society. This study has highlighted the importance of specific datasets and quantitative methods to investigate Dutch culinary trends, prompting the need for tailored methodologies to address domain-specific historical texts. The potential of LLMs to address this need has proved particularly promising, paving the way for future directions in this research.

Author Contributions (by author initials) are listed according to the Contributor Roles Taxonomy (CRediT).[4] Conceptualization: GB, MvE, TP. Data curation: GB Writing (original draft): GB, MvE, RH. Writing (review and editing): GB, MvE, TP, RH.

Acknowledgements
This research was funded by the European Union under grant agreement 101088548 - TRIFECTA. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. The authors thank Thijs Elfrink for his insightful and valuable feedback on the final draft.

References

Jenny Arendholz, Wolfram Bublitz, Monika Kirner-Ludwig, and Iris Zimmermann. 2013. Food for thought–or, what's (in) a recipe? a diachronic analysis of cooking instructions. In *Culinary Linguistics: The chef's special*, pages 119–138. John Benjamins Publishing Company.

Peter Berg and Ruth Ann Jones. 2003. Feeding america: The historic american cookbook project. *Journal of Agricultural & Food Information*, 5(4):69–75.

Usashi Chatterjee, Fausto Giunchiglia, Devika P Madalli, and Vincenzo Maltese. 2016. Modeling recipes for online search. In *OTM 2016 Conferences, Rhodes, Greece, October 24-28, 2016, Proceedings*, pages 625–642. Springer.

Kyri W Claflin. 2020. What shall we have today?: Culinary adaptation in world war i paris. In *Dublin Gastronomy Symposium*.

Fred Czarra. 2009. *Spices: a global history*. Reaktion Books.

David Elsweiler, Christoph Trattner, and Morgan Harvey. 2017. Exploiting food choice biases for healthier recipe recommendation. In *Proceedings of the 40th international acm sigir conference on research and development in information retrieval*, pages 575–584.

Steven Haussmann, Oshani Seneviratne, Yu Chen, Yarden Ne'eman, James Codella, Ching-Hua Chen, Deborah L McGuinness, and Mohammed J Zaki. 2019. Foodkg: a semantics-driven knowledge graph for food recommendation. In *The Semantic Web–ISWC 2019: 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part II 18*, pages 146–162. Springer.

Anjali Jain and Alka Singhal. 2022. Personalized food recommendation—state of art and review. *Ambient communications and computer systems: Proceedings of RACCCS 2021*, pages 153–164.

Elaine Leong. 2019. *Recipes and everyday knowledge: Medicine, science, and the household in early modern England*. University of Chicago Press.

Javier Marın, Aritro Biswas, Ferda Ofli, Nicholas Hynes, Amaia Salvador, Yusuf Aytar, Ingmar Weber, and Antonio Torralba. 2021. Recipe1m+: A dataset for learning cross-modal embeddings for cooking recipes and food images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):187–203.

Floor Meijer. 2014. Vrouwenlexicon van nederland, wannée, cornelia johanna (1880-1932).

Susie Protschky. 2008. The colonial table: Food, culture and dutch identity in colonial indonesia. *Australian Journal of Politics & History*, 54(3):346–357.

Sola S Shirai, Oshani Seneviratne, Minor E Gordon, Ching-Hua Chen, and Deborah L McGuinness. 2021. Identifying ingredient substitutions using a knowledge graph of food. *Frontiers in Artificial Intelligence*, 3:621766.

Albert Sonnenfeld, Jean-Louis Flandrin, and Massimo Montanari. 1999. *Food: a culinary history*. Columbia University Press.

Chun-Yuen Teng, Yu-Ru Lin, and Lada A Adamic. 2012. Recipe recommendation using ingredient networks. In *Proceedings of the 4th annual ACM web science conference*, pages 298–307.

M. van Erp, M. Wevers, and H. Huurdeman. 2018. Constructing a recipe web from historical newspapers. In D. Vrandečić et al., editors, *The Semantic Web – ISWC 2018*, volume 11136 of *Lecture Notes in Computer Science*, pages 189–206. Springer.

H. Verwey-Jonker. 1955. Social services in the netherlands. *Marriage and Family Living*, 17(3):233–235.

[4] https://credit.niso.org/

A Appendix

Legend for LLM Parsing Result

Table 2 in the article gives the results of the parsing of recipes with LLM. The recipes were first parsed to a json representation using the mistral LLM model. The json results of the LLM parsing were then converted to an xml representation, (1) with a rule based json-xml conversion; (2) with an LLM-based json-xml conversion. This is more complicated than a straightforward comparison of the json output, but allows for a more structured comparison of the original recipe and the output. Table 2 indicates different statuses for the final output, comparing the rule-based XML and the LLM-based XML, evaluating for the existence of an <ingredienten> (ingredients) and a <bereiding> (preparation) section. There are the following possible outcomes:

- "original_more_complete": Indicates that the original XML output has both "<ingredienten>" and "<bereiding>" tags, while the LLM-generated output is missing one or both of these key elements. "details": Explains that the original XML is more complete. "better": "original", as the original XML contains all key elements.
- "llm_more_complete": Indicates that the LLM-generated XML output has both "<ingredienten>" and "<bereiding>" tags, while the original XML output is missing one or both of these key elements. "details": Explains that the LLM XML is more complete. "better": "llm", as the LLM XML contains all key elements.
- "original_more_content": Indicates that the original XML output has significantly more text content than the LLM-generated output (difference in length > 50 characters after stripping tags). "details": Explains that the original XML has significantly more content. "better": "original", as the original XML contains more information.
- "llm_more_content": Indicates that the LLM-generated XML output has significantly more text content than the original XML output (difference in length > 50 characters after stripping tags). "details": Explains that the LLM XML has significantly more content. "better": "llm", as the LLM XML contains more information.

Missing words	Proportion
0-1	44%
1-5	25%
5-10	16%
10+	14%

Table 4: Comparison of missing words between the recipe text and the LLM

- "similar": Indicates that both XML outputs are structurally similar and have comparable content. This is determined by passing all previous checks, including error checks, completeness, and content length comparison. "details": Explains that both XML outputs are structurally similar and have comparable content. "better": "both", as both outputs are considered equally good.
- "parsing_error": Indicates an error occurred in parsing either the original or LLM-generated XML using BeautifulSoup. "details": Provides the specific error message encountered during parsing. "better": "neither", as a parsing error prevents proper evaluation.

For reasons of structure there is no direct comparison with the source recipe text, as the absence of structure makes this complicated. Table 4 shows a (shallow) comparison of missing words between the recipe text and the LLM JSON results.

Translations for Dutch ingredients mentioned in the heatmap visualisation

Dutch Ingredient	English Translation
Aardappel	Potatoes
Bloem	Flour
Boter	Butter
Bouillon	Broth
Brood	Bread
Citroen	Lemon
Ei	Egg
Eiwit	Egg White
Kaas	Cheese
Melk	Milk
Margarine	Margarine
Olie	Oil
Peper	Pepper
Peterselie	Parsley
Room	Cream
Saus	Sauce
Suiker	Sugar
Ui	Onion
Vanille	Vanilla
Zout	Salt

Table 5: Translation of Dutch Ingredients from Fig 2