Eurographics Conference on Visualization (EuroVis) 2016
E. Bertini, N. Elmqvist, and T. Wischgoll
(Guest Editors)

Short Paper

Classic Techniques in New Domains: An Alternative Recipe

M. Monroe

IBM, Cambridge, MA

Abstract

In this paper, we adapt the classic technique of depicting a process as a structured workflow to suit the standard recipe. Cooking can be thought of as a small data, big user task. A single recipe encompasses only a small amount of information, but is utilized across a large user base. Our goal was to understand and measure the benefits of tailoring the presentation of a recipe to suit a specific faction of users. As such, our more technical rendering was paired with a technically proficient user base, resulting in dramatic gains in both the speed and accuracy with which the information was interpreted. These benefits serve to motivate our continued work towards automatically translating recipes into a structured data format that can be easily reconfigured into this and other representations of the information to enable a more customized experience across a large and varied user base.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [INFORMATION INTERFACES AND PRESENTATION]: User Interfaces—User-centered design

1. Introduction

When compared to the current focus of most visualization and analytics research, cooking is decidedly a “small data" task. Even with the most lavish augmentations, it is hard to imagine a single recipe exceeding a few megabytes of information. However, it is just as decidedly a “big user" task. That is, it is a task that effects every human on earth (either directly or indirectly) multiple times a day. As such, improvements to the way that this task is performed, even if they affect only a subset of the total user base, could have a significant overall impact. In this paper, we investigate the benefits of depicting a recipe as a structured workflow. This format was selected for two key reasons:

1. It addresses many of the difficulties of using the standard recipe format that were brought up during our initial user surveys.
2. It stems from a readily apparent, underlying data structure, meaning that it could conceivably be realized using automated parsing techniques.

While the latter capability remains the focus of continued development, this paper is intended serve as the proverbial “carrot on a stick," demonstrating that the workflow format can have a significant impact on the speed and accuracy with which recipes are interpreted. Our technically inclined users found the design to be intuitive and practical across multiple proposed scenarios. Finally, we constructed a prototype recipe editor that can help to bridge the gap between manual and automated conversions of the standard recipe format into the workflow format.

2. Background

The current, standard format of the recipe revolves around a natural language, written description of the cooking process. The recipe’s ingredients, their measurements, and any “prep" work that needs to be done to them (chopping, mincing, grating, etc.) are typically listed in a separate section, either above the written description [Web15a, Web15d] or to the left of it [Web15c, Web15b]. The total time that is required to execute the recipe is listed as well, often as the additive value of the “prep time" and the “cooking time."

Perhaps the most interesting thing about the evolution of this standard recipe format over the past century is that it is virtually nonexistent. Figure 1 depicts a recipe from the start of the twentieth century next to a recipe from a typical website today. They are functionally identical. This is not to say that innovation has been entirely lacking. The design world has envisioned a series of drastically inventive recipe formats (see Figure 2) [Twe11], and various cookbooks and mobile apps have dreamed up unique presentations [Sid14, Fer12]. However, translating a recipe from the standard format into these more creative formats would be a predominately manual and time-consuming process.

On the computational end of the spectrum, automated recipe parsing has been attempted in a number of previous works [TS08,MWCM14,AM08,TLA12], including work that has focused specifically on extracting the workflow of the cooking process [MBGW10,KW11]. However, this work revolves around building a complete semantic model that might allow for algorithmic reasoning or robotic execution. The potential of leveraging these parsing capabilities to create custom experiences for human users has yet to be explored.

DOI: 10.2312/eurovisshort.20161172

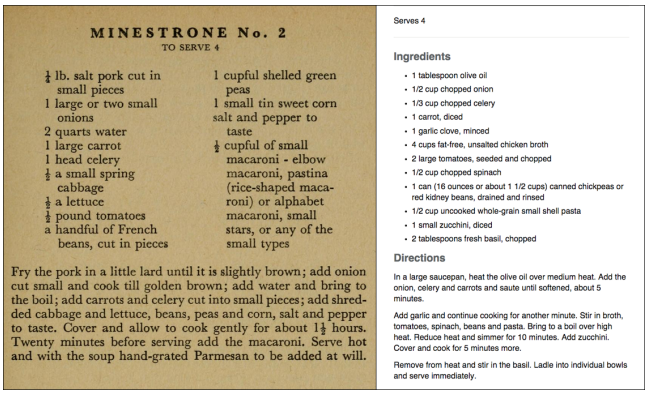

MINESTRONE No. 2

TO SERVE 4

¼ lb. salt pork cut in small pieces
1 large or two small onions
2 quarts water
1 large carrot
1 head celery
½ a small spring cabbage
½ a lettuce
½ pound tomatoes
a handful of French beans, cut in pieces
1 cupful shelled green peas
1 small tin sweet corn
salt and pepper to taste
½ cupful of small macaroni - elbow macaroni, pastina (rice-shaped macaroni) or alphabet macaroni, small stars, or any of the small types

Fry the pork in a little lard until it is slightly brown; add onion cut small and cook till golden brown; add water and bring to the boil; add carrots and celery cut into small pieces; add shredded cabbage and lettuce, beans, peas and corn, salt and pepper to taste. Cover and allow to cook gently for about 1½ hours. Twenty minutes before serving add the macaroni. Serve hot and with the soup hand-grated Parmesan to be added at will.

Serves 4

Ingredients

- 1 tablespoon olive oil
- 1/2 cup chopped onion
- 1/3 cup chopped celery
- 1 carrot, diced
- 1 garlic clove, minced
- 4 cups fat-free, unsalted chicken broth
- 2 large tomatoes, seeded and chopped
- 1/2 cup chopped spinach
- 1 can (16 ounces or about 1 1/2 cups) canned chickpeas or red kidney beans, drained and rinsed
- 1/2 cup uncooked whole-grain small shell pasta
- 1 small zucchini, diced
- 2 tablespoons fresh basil, chopped

Directions

In a large saucepan, heat the olive oil over medium heat. Add the onion, celery and carrots and saute until softened, about 5 minutes.

Add garlic and continue cooking for another minute. Stir in broth, tomatoes, spinach, beans and pasta. Bring to a boil over high heat. Reduce heat and simmer for 10 minutes. Add zucchini. Cover and cook for 5 minutes more.

Remove from heat and stir in the basil. Ladle into individual bowls and serve immediately.

Figure 1: *Two recipes for minestrone soup, one from 1900 (left) [Dal00], and one from a current website (right) [Rec15b]. The two recipe formats are more or less equivalent.*

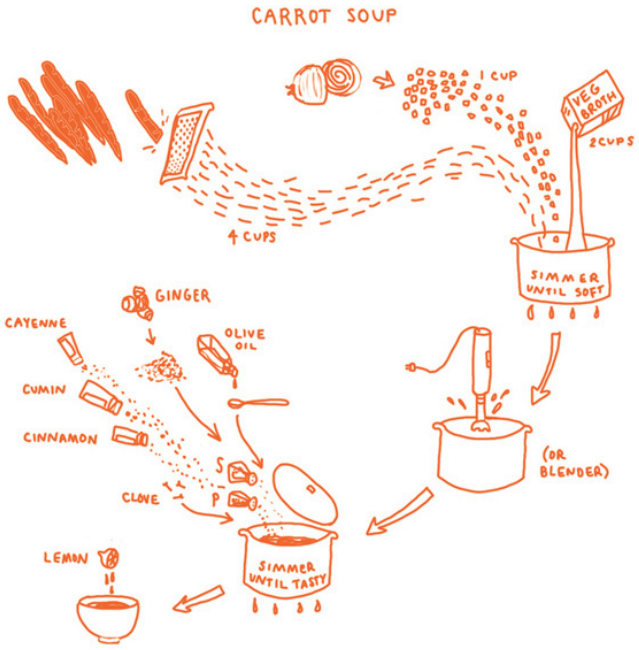

Figure 2: *Visual solutions from professional designers can be extremely creative [She13], but difficult to replicate at scale.*

3. Catering to the Technically Inclined

In order to understand the benefits of tailoring the recipe to suit a particular subgroup of users, we conducted a user study comprised of 20 technically inclined participants. These participants, 11 male and 9 female, were recruited from a lab of scientific researchers and developers. Apart from their consistent, technical backgrounds, the participants ranged from novice to experienced cooks, and spanned an age range of over 30 years (youngest participants in their mid-twenties, oldest participants over sixty years old).

We asked our 20 participants to describe their biggest challenges when following the standard recipe format. Across their responses, the most common answer was that it is extremely difficult to look at a recipe and quickly determine exactly what they are getting themselves into. Without carefully reading the entire cooking process and ingredient list, it is hard to determine:

- How much of the total execution time involves hands-on engagement.
- Whether the recipe includes specialty or hard to find ingredients or equipment.
- The total number of kitchen resources required (pots, pans, burners, ovens).
- Whether the recipe involves any difficult cooking techniques.
- Whether the total execution time is accurate.

Furthermore, the standard recipe format is designed to support a cooking process that is comprised of two mutually exclusive steps: preparing the ingredients and cooking the dish. This strategy plays a critical role in professional kitchens and what is known as *mise en place*. Professional kitchens are able to produce finished meals very quickly because they prepare their ingredients "off the clock." That is to say, ingredients are prepared before the restaurant opens.

But for home cooks, there is no such thing as "off the clock," and as a result, many users attempt to interleave these two steps. Of our 20 study participants, only 8 said that they prep all of their ingredients before starting to cook. While this strategy can significantly reduce the time it takes to complete a recipe, it creates the added challenge of ensuring that ingredients are ready when the recipe calls for them.

3.1. Workflow Redesign

To address these challenges, we created an alternate recipe format, intended to enable a more systematic approach to the cooking process. The design drew heavily on foundational visualization strategies, such as workflow diagrams and Gantt charts [Tuf83, GG21, Gan13], but was customized specifically for the cooking domain.

Our foremost focus was to represent the timing of the recipe as unambiguously as possible. To do this, we plotted each step of the recipe against a timeline, as shown in Figure 3. Any step that required a direct action from the user - a "hand-on" step - was represented as black ellipse containing an instruction such as "Sauté" or "Bring to boil." Instructions that included the addition of an ingredient were represented in green, with the ingredient(s) and their measurements listed to the side of the instruction. When an ingredient required more preparation than a simple measurement (chopping, mincing, grating, etc.), it was listed in a dark green tag above the ingredient. Any down time or waiting time (such as waiting for water to boil) was represented as a gray extension to the previous instruction.

The final component of this workflow design was that multiple timelines could be represented in parallel. This represented tasks that could be executed independently, using a unique set of equipment, such as cooking pasta in a pot while separately preparing a sauce in a pan. A timeline could come to an end if its contents were added to a different timeline. However, since a single chef can typically perform only one task at a time, the "hands-on" instructions in each timeline were not allowed to overlap.

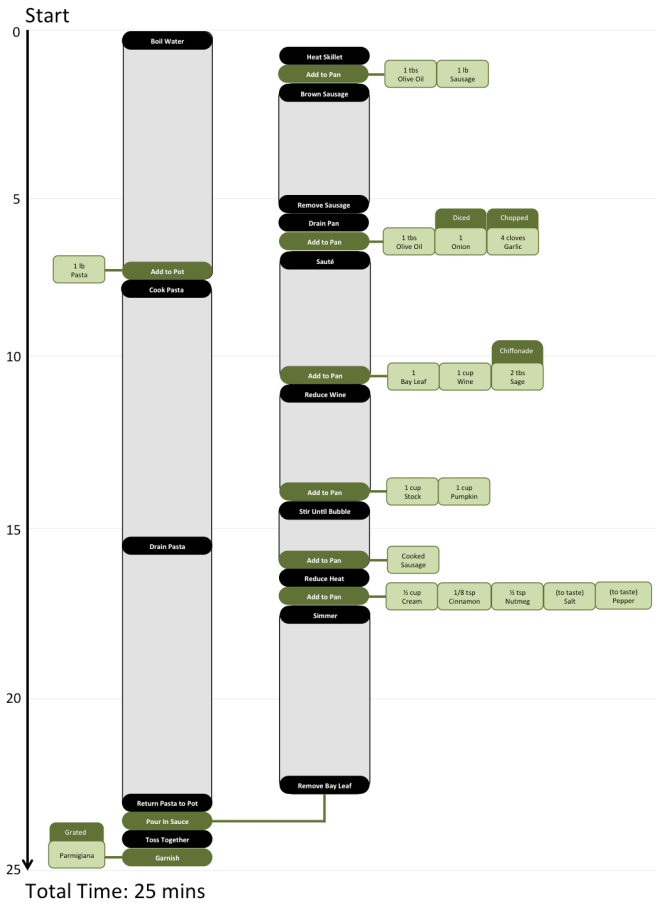

Figure 3: *Our workflow redesign presents a recipe as a timeline of staggered tasks.*

4. User Study

We asked our 20 study participants to answer a series of 6 questions about two similar recipes [Rec15c, Rec15a], one displayed in the standard format and one displayed in the workflow format described above. The questions were designed to reflect the challenges described in Section 3. Since our primary objective was to observe the participants' comprehension of the information, rather than their functional cooking skills, participants were not asked to actually execute the recipe. However, to simulate the task switching of an actual cooking experience, in which the cook must constantly shift their gaze from the recipe to the task at hand, participants were required to turn away from the recipe (displayed on a 27" monitor) while each of the following questions was being asked:

Q1: Which ingredients need to be prepped before being added into this recipe?

Q2: You've just added the cooked sausage back into the pan and reduced heat, what ingredients must you add next and in what quantities?

Q3: You've just drained the pasta, and a friend walks in and asks what he/she can do to help. What's your suggestion?

Q4: You just returned the pasta to the cooking pot. What's your next step?

Q5: You need about three minutes to set the table. When would be a good time to do that?

Q6: You want to call out a "10-minute warning" to the rest of the family, when should you do that?

Participants received a brief tutorial on the recipe format that they would be using first, during which they could ask questions about the formatting. They were then shown the recipe that they would be "cooking" and given time to look the recipe over, much like they would presumably do if they were preparing to cook the dish. The participants were then asked the 6 questions listed above and their answers were recorded and scored for both accuracy (correct or incorrect) and the time it took to produce them. Because the cooking process can be highly personalized, we tried to be as lenient as possible when coding the participants' responses for accuracy. The study was counterbalanced to control for order effects, both in terms of the recipe that each participant saw first, as well as the recipe format.

5. Results

Using our workflow recipe design, participants produced nearly double the number of correct answers than they produced using the standard recipe format. Additionally, they were able to produce these answers in half the time that they took using the standard format. A breakdown of the accuracy and response times by question can be found in Figure 4. Perhaps more telling than these quantitative results, however, are the strategies that produced them:

Strategic Preparation: Before a single question was asked, participants using the workflow design appeared to arrive at a higher level understanding of the recipe execution than they did using the standard format. One participant mentioned that he would probably start by combining five of the ingredients into a single bowl since they would eventually be added together, a conclusion he reached after only a few seconds of scanning the recipe. Another participant quickly noted that one section of the recipe was particularly "prep-heavy," an observation that would have been difficult to arrive at using the standard format.

Q1: To answer Q1 using the standard recipe format, participants had no choice but to read the entire ingredient list, looking for the action verbs that indicate whether prep work is needed. Some participants read portions of the written description as well, even though they were told upfront that prep requirements would be listed in the ingredient list. By contrast, participants using the workflow design could quickly scan the display for the green prep tags, a process that was significantly faster to execute.

Q2 and Q4: Q2 and Q4 targeted the most fundamental question that cooks must ask when following a recipe: What do I do next? Of the 6 questions that we posed to our user study participants, these two questions had the most clear-cut, ground truth

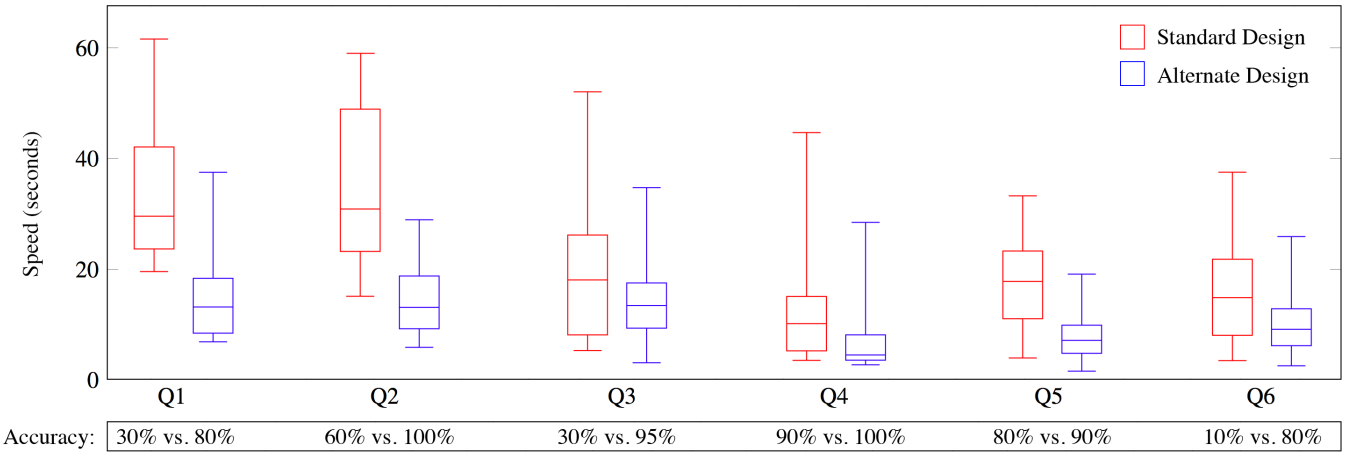

Figure 4: *Using the workflow design, our study participants produced significantly faster answers across all 6 questions ($p < .01$ using the Wilcoxon Signed-Rank Test). Additionally, they consistently produced more accurate answers using the workflow design.*

answers and were the least likely to be influenced by individual cooking strategies or skillsets. Still, participants using the standard recipe format were significantly slower and less accurate than they were using the workflow design. Q2 was the most telling. Using the standard format, participants not only had to find the correct place in the recipe text, but then had to cross reference the prescribed ingredients with the measurement quantities in the ingredient list. Many participants resorted to placing a finger on the recipe text to ensure that they did not lose their place in the written description while they scanned the ingredient list, a strategy that would not be as easy to employ in an actual kitchen. Using the workflow design, participants were able to find their place in the recipe faster and read the ingredients and their quantities inline.

Q3: This question gave participants the most leeway in terms of personal preference and strategy. However, it also relied heavily on a high-level understanding of how the recipe needed to be executed. This proved to be extremely difficult using the standard recipe format. Many users were so stumped by this question that many blurted out a facetious answer (“Pour me a glass of wine!") before begrudgingly searching for a more reasonable one. Using the workflow design, however, participants produced a wide array of detailed responses about how to optimize their efforts.

Q5 and Q6: Participants appeared to employ a consistent strategy for answering these two questions using the standard recipe format. They would scan the recipe text for numbers, looking for explicitly stated time lapses. This strategy was effective for Q5, since both recipes included at least one explicitly stated time lapse of the appropriate size. However, for Q6, this strategy was disastrous. Participants tried to estimate a ten-minute period of time by scanning for explicitly stated time lapses and adding them together until they hit ten minutes. In doing this, however, they missed major time lapses that were not described numerically, such as “bring to boil" or “reduce by half." These time lapses were presumably also missed when participants looked for down time during which they could prep ingredients.

Overall Impression: Our participants were overwhelmingly positive about the workflow design, stating either stating an overt preference or listing series of advantages that it had over the standard format. Even participants with extensive cooking experience, who stated that they enjoyed the natural language, story-like component of the standard recipe format, admitted that certain questions were much easier to answer using the workflow design. Many participants suggested additional settings in which the workflow design would be useful, such as cooking with a partner, or cooking multiple recipes concurrently.

6. Discussion and Future Work

In this paper, we demonstrated that a more technical presentation of a recipe can result in significant speed and accuracy gains across a technically inclined user base. This motivates our continued work in this domain to enable the automatic conversion of recipes from the standard format into our proposed workflow format. To help bridge the gap between manual and automatic generation of our workflow design, we constructed a web-based recipe editor. Users click a canvas to create instruction ellipses, and can then drag the edges of these ellipses to add “hands-on" duration, “down-time" duration, or ingredient additions. These instructions can then be edited and positioned as needed on the canvas.

This editor not only allows users to construct recipes from scratch, but can also serve as the output for automated parsing approaches, which can then be manually fine-tuned. For example, determining how to break a recipe into parallel timelines is extremely difficult for current automated approaches. However, if an automated parse attempt could create a single timeline, users could drag the appropriate sections of it into separate timelines. Current annotation efforts such as the MILK language [TS08] will help us to determine how best to manage these tradeoffs. Ultimately, the goal is to arrive at an underlying representation of the recipe that can be displayed not only in our proposed workflow format, but in an array of formats that suit different factions of this large and varied user base.

References

[AM08] AGARWAL R., MILLER K.: *Information Extraction from Recipes*. Tech. Rep. CS224N-Ling284-AgarwalMiller, Department of Computer Science, Stanford University, Stanford, CA, 2008. 1

[Dal00] DALY D.: *Italian Cooking*, first ed. Spring Books, 1900. 2

[Fer12] FERRIS T.: *The 4-Hour Chef*, first ed. New Harvest, 2012. 1

[Gan13] GANTT H. L.: Work, wages, and profits. 2

[GG21] GILBRETH F., GILBRETH L. M.: Process charts - first steps in finding the one best way. 2

[KW11] KIRSTIN WALTER MIRJAM MINOR R. B.: Workflow extraction from cooking recipes. In *Workshop Proceedings ICCBR 2011* (2011), pp. 207–216. 1

[MBGW10] MINOR M., BERGMANN R., GÖRG S., WALTER K.: Adaptation of cooking instructions following the workflow paradigm. In *Workshop Proceedings ICCBR 2010* (2010), pp. 199–208. 1

[MWCM14] MALMAUD J., WAGNER E. J., CHANG N., MURPHY K.: Cooking with semantics. In *ACL 2014 Workshop on Semantic Parsing (SP14)* (2014). 1

[Rec15a] RECIPE: CHRISTMAS PASTA: http://www.foodnetwork.com/recipes/rachael-ray/christmas-pasta-recipe2.html. 2015. Accessed: 2015-02-19. 3

[Rec15b] RECIPE: MINESTRONE SOUP: http://www.mayoclinic.org/healthy-living/recipes/minestrone-soup/rcp-20049680. 2015. Accessed: 2015-02-13. 2

[Rec15c] RECIPE: PASTA WITH PUMPKIN AND SAUSAGE: http://www.foodnetwork.com/recipes/rachael-ray/pasta-with-pumpkin-and-sausage-recipe.html. 2015. Accessed: 2015-02-19. 3

[She13] SHELLY K.: *Picture Cook*, first ed. Ulysses Press, 2013. 2

[Sid14] SIDECHEF INC.: SideChef. Apple and Android App Stores, 2014. 1

[TLA12] TENG C.-Y., LIN Y.-R., ADAMIC L. A.: Recipe recommendation using ingredient networks. In *Proceedings of the 3rd Annual ACM Web Science Conference* (2012), pp. 298–307. 1

[TS08] TASSE D., SMITH N. A.: *SOUR CREAM: Toward Semantic Processing of Recipes*. Tech. Rep. CMU-LTI-08-005, Language Technologies Institute, Carnegie Mellon University, Pittsburgh, PA, 2008. 1, 4

[Tuf83] TUFTE E. R.: The visual display of quantitative information. 2

[Twe11] TWENTY-TWO MIND-BLOWING SUBMISSIONS FOR OUR REDESIGN THE RECIPE PROJECT SUBMISSIONS: REDESIGN THE RECIPE: http://magazine.good.is/slideshows/submission-redesign-the-recipe. Good Magazine, 2011. Accessed: 2015-02-13. 1

[Web15a] WEBSITE: ALLRECIPES.COM: http://allrecipes.com. 2015. Accessed: 2015-02-13. 1

[Web15b] WEBSITE: FOOD.COM: http://www.food.com. 2015. Accessed: 2015-02-13. 1

[Web15c] WEBSITE: FOODNETWORK.COM: http://www.foodnetwork.com. 2015. Accessed: 2015-02-13. 1

[Web15d] WEBSITE: RECIPE.COM: http://www.recipe.com. 2015. Accessed: 2015-02-13. 1