Rapid Prototyping of Mobile Apps for Clinical Research using Semantic Web Technologies *

Oshani Seneviratne[1], William Van Woensel[2], Giuseppe Loseto[3], Floriano Scioscia[3], Evan W. Patton[4], and Lalana Kagal[4]

[1] Rensselaer Polytechnic Institute, Troy NY 12180, USA
senevo@rpi.edu
[2] Dalhousie University, Halifax NS B3H 4R2, Canada
william.van.woensel@dal.ca
[3] Polytechnic University of Bari, 70126 Bari BA, Italy
giuseppe.loseto@poliba.it, floriano.scioscia@poliba.it
[4] Massachusetts Institute of Technology, Cambridge MA 02143, USA
ewpatton@mit.edu, lkagal@mit.edu

Abstract. This paper provides a demo of the Punya framework, introduced in our resource track paper titled "The Punya Platform: Building Mobile Research Apps with Linked Data and Semantic Features." In this demo, we focus on using the Punya framework to build clinical research apps with a particular focus on Type 2 diabetes research. We demonstrate how Semantic Web and biomedical researchers can utilize Punya to rapidly prototype mobile semantic apps. Apps developed with Punya can consume Linked Data, thus exploiting the vast biomedical data sources available on the Web; while producing RDF from the app for downstream analysis. Researchers can also write semantic rules to encode clinical guideline recommendations in a visual manner within the Punya framework. Due to the intuitive and visual editing environment, researchers can easily tweak complex decision-logic rules.

Submission Type: Demo
Link: https://punya.mit.edu/#use-cases

Keywords: Mobile App Development · Research Apps · Clinical Apps · Linked Data · Semantic Rules · Rapid Prototyping

1 Introduction

Punya is a Semantic Web-enabled app development framework built on top of MIT App Inventor [9]. The user-friendly, block-based programming paradigm provided in MIT App Inventor allows even non-programmers to quickly and easily create Android mobile apps. Specifically, the designer view allows users to drag-and-drop visual objects to create a user interface, whereas the block view allows users to "piece" together application logic using control, text, math, and other operators as puzzle pieces. Punya adds semantic capabilities to MIT App Inventor, including reading and writing Linked Data, semantic reasoning, and connecting to sensors using both low-level Bluetooth APIs and high-level Linked Data Platform over Constrained Application Protocol (LDP-CoAP) [5]. Advanced capabilities include an embedded semantic rule engine, visual editor for authoring semantic rules, and integrating GraphQL and SPARQL to access and query remote graph data. Leveraging Linked Data enables developers to reuse the thousands of valuable datasets from the Linked Open Data cloud in their mobile apps; while the generated Linked Data can be analyzed, combined and integrated with other Linked Data (e.g., generated by other apps).

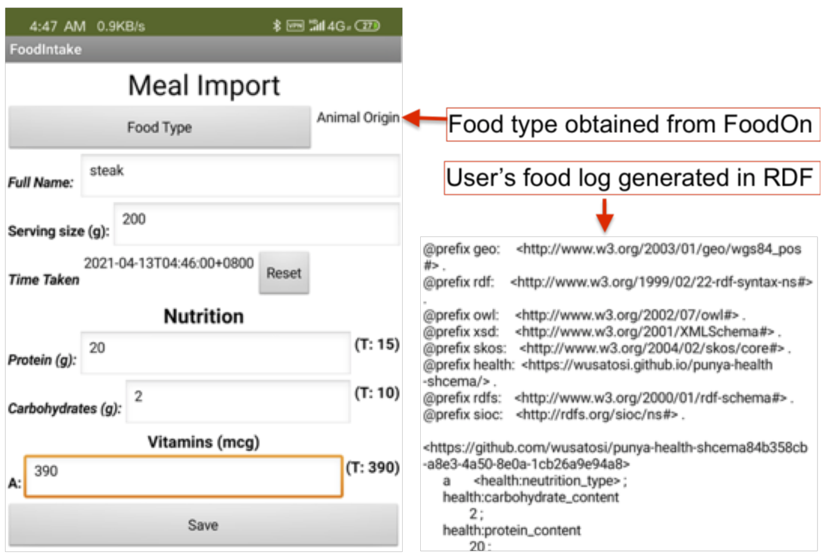

Fig. 1: Example meal input application for diabetes management containing semantically annotated input fields producing an RDF representation of the food log. Metadata about the food item will be obtained from external data sources (i.e., food type from FoodOn [2], and the nutrition information from the FoodKG [3]).

As smartphones have become pervasive, many clinical researchers are looking to utilize patients' mobile context to improve the self-management of diseases and gain behavioral insights. Semantic Web technologies have been widely adopted in healthcare research to support these efforts [10]. Moreover, many ontologies and Semantic Web resources are usable off-the-shelf. Using Punya, researchers can rapidly prototype mobile semantic apps that consume such semantic biomedical data, implement local decision logic, and produce RDF data directly from the app itself. We note that the generated RDF data further contributes to the Linked Data Cloud and may thus provide an answer to the perennial chicken-and-egg problem on data generation in the Semantic Web [4].

2 Clinical Research Application Development with Punya

As an example use case for Punya in action, we illustrate how a clinical researcher working with Type 2 Diabetes (T2D) patients can quickly prototype a mobile app. T2D is a health condition affecting a significant proportion of the world population [8,7]. The importance of behavior change in self-managing T2D is well recognized, and a variety of psychological theories may be applied (e.g., Social Cognitive Theory, Cognitive Behavioral Therapy). Moreover, context-specific interventions can be offered on smartphones in many different ways (e.g., using different modalities). These factors point towards the need for rapid prototyping of a variety of mobile self-management approaches, and evaluating their effectiveness. Furthermore, given that there is no one-size-fits-all solution for behavior modification, many clinical researchers have turned to personalized apps, which effectively capture patient data and relay any notifications as per the patients' needs and context.

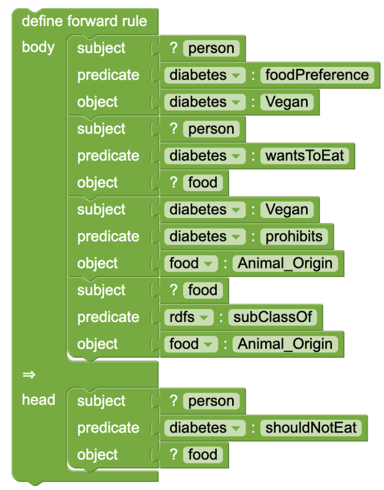

Fig. 2: A semantic rule that determines if a given food item is suitable for a vegan individual based on the food's classification available in FoodOn [2].

Using nutrition behavior as a focus, we demonstrate how existing ontologies such as FoodOn [2] can be used to capture a patient's daily food intake using a mobile app developed using Punya as shown in Fig. 1. The screen on the left-hand side shows the user interface where a user can input the meal name and the serving size. Once the information is input to the app, the nutritional content is auto-populated from external sources such as the FoodKG [3]. The right-hand side image in Fig. 1 shows the RDF data of the user's food log that is generated when the form is submitted. This RDF food log may be saved on the phone or submitted to an external data endpoint. As can be seen in the figure, Punya allows querying external data sources: when the user enters "steak," the app queries FoodOn and ascertains it is a food type of *Animal Origin*. By exploring the food type hierarchy in FoodOn for a given food item, we can identify a substitution for *steak* in case the user is sharing the meal with someone who is maintaining a particular type of diet, such as a vegan diet as they cannot consume a food type of *Animal Origin*. Understanding and capturing context this way, with data augmented by the Linked Data cloud, can provide additional benefits. Extending the previous idea, Fig. 2 demonstrates a rule set that can be used to determine whether a vegan user can consume a given food item. The rule could be further extended to determine a substitution food category for the user.

Additionally, as some diabetic patients are required to take insulin, there is a need to determine the correct insulin dosage based on specific parameters of the patient, such as blood glucose levels, height, weight, etc. At the same time, certain foods high in a particular nutrient may have a contraindication with the drug being taken. For example, if a researcher is studying whether *protein* in a diabetic patient has a contraindication with *insulin* administered, the researcher could utilize a rule such as the one give given in Fig. 3. The visual semantic rule could easily be tweaked with the threshold for a specific patient population studied. By using the food input from Fig. 1, and the auto-populated protein value from the FoodKG [3], the Punya semantic reasoner can provide a recommendation as to whether this particular food item is suitable or not for the user. This visual format makes rule authoring much more accessible to clinical researchers, who may not be well-versed in Semantic Web technologies such as RDF and OWL. Furthermore, the output from this rule could be wired into the UI of Fig. 1 to alert the user if needed.

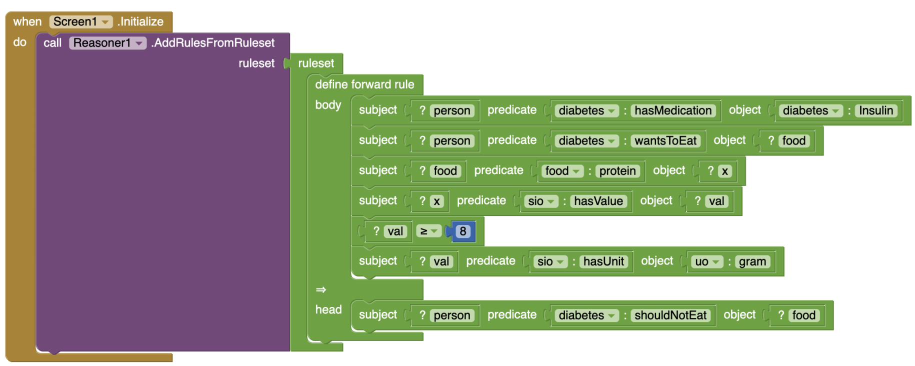

Fig. 3: A semantic rule encoding that shows that a diabetic person should not consume a food item if the protein value exceeds a certain threshold. The rule is input into the *ruleset* of the *Reasoner* component, and the values in the rule can be tweaked easily.

3 Related Work

Apple provides a software framework called ResearchKit [1] that lets medical researchers develop surveys, forms, and activities to gather study data. However, it is beyond its scope to query online health data sources, implement complex application logic, or support semantic technologies. In contrast, the Punya platform provides a convenient drag-and-drop environment for researchers to quickly prototype their ideas and easily access Linked Data resources. Node-RED [6] provides a browser-based editor to wire up event-based Internet of Things (IoT) applications quickly, with a similar drag-and-drop interface. However, Node-RED focuses on setting up data flows, including IoT devices and online services, while Punya focuses on prototyping mobile (research) apps with a fully-fledged UI and Semantic Web support.

4 Conclusion

This paper demonstrates how mobile research apps in the clinical space could benefit from Punya, with a simple meal input and clinical reasoning example. Other features in Punya, such as LDP-CoAP, can be used to quantify patient behaviors (e.g., medication adherence) more objectively by relying on pervasive sensors (e.g., smart pillbox) rather than inaccurate personal reporting. In conclusion, Punya provides an easy-to-use framework for developing mobile (research) apps in many domains. We have included comprehensive documentation with Punya and have provided sample apps in our online resource. The Punya framework also includes built-in tutorials to assist anyone who wants to get started with rapid app development.

Acknowledgements

This work is partially supported by IBM Research AI through the AI Horizons Network, and University of New Brunswick Research Fund Competition 2020 (RF Explore).

References

1. Apple Inc.: ResearchKit, http://researchkit.org/
2. Dooley et al.: Foodon: a harmonized food ontology to increase global food traceability, quality control and data integration. NPJ Science of Food **2**(1), 1–10 (2018)
3. Haussmann, S., Seneviratne, O., Chen, Y., Ne'eman, Y., Codella, J., Chen, C.H., McGuinness, D.L., Zaki, M.J.: Foodkg: a semantics-driven knowledge graph for food recommendation. In: International Semantic Web Conference. pp. 146–162. Springer (2019)
4. Hogan, A.: The semantic web: Two decades on. Semantic Web **11**(1), 169–185 (2020)
5. Loseto, G., Ieva, S., Gramegna, F., Ruta, M., Scioscia, F., Di Sciascio, E.: Linked Data (in low-resource) Platforms: a mapping for Constrained Application Protocol. In: International Semantic Web Conference. pp. 131–139. Springer (2016)
6. Node-RED community: Node-RED: Low-code programming for event-driven applications, https://nodered.org
7. Sharma, N.C.: Government survey found 11.8% prevalence of diabetes in india. LiveMint (2019)
8. US CDC: National diabetes statistics report (2020)
9. Wolber, D., Abelson, H., Friedman, M.: Democratizing computing with app inventor. GetMobile: Mobile Computing and Communications **18**(4), 53–58 (2015)
10. Zenuni, X., Raufi, B., Ismaili, F., Ajdari, J.: State of the Art of Semantic Web for Healthcare. Procedia-Social and Behavioral Sciences **195**, 1990–1998 (2015)