Auto-ID Labs ETH/HSG

# Year in Review

2020

# Content




Date of report: Sep 2020

Cover photo credit: Pinterest

# Director’s note

Figure 1: Prof. Elgar Fleisch (left) and research team at ETH Zurich

Dear colleagues and friends,

On behalf of all our lab members, I'm delighted to share with you a few highlights of the research conducted at the Swiss Auto-ID Labs in the past year.

Topics that we started a few years ago received now significant traction in the industry. For example, the digital receipts topic is discussed globally not only as a green alternative to paper receipts, but also as enabler for new business applications. We are very excited that the research resulted in a Special Interest Group of GS1 and pushing this now to the next level towards standardization. Our universities, HSG and ETH, both were awarded with research grants for continuing the work on digital receipts and future applications of the new data in the domains of health and sustainability. Hence, our work in raising nutritional awareness & leveraging product databases has gained significant traction in Europe and might contribute to healthier shopping habits.

Driven by the global advancements in Computer Vision & Artificial Intelligence, we also started again to develop new areas that rely on visual object and context recognition. Not, because the barcode-based product recognition is obsolete, but because mixed reality headsets and computer vision enabled applications promise revolutionary new use-cases, ranging from augmented reality promotion, to real-time store-shelf monitoring and cashier-less checkouts. The insights of this work can become important for the GS1 community to understand how to navigate the Augmented Reality area and to inform about the coexistence of barcodes and object-recognition and scenarios that fully rely on visual recognition. As part of learning based systems, we have also started a first project to bridge the gap between artificial intelligence systems and the relation to humans. We started a project in the manufacturing execution system area that can provide a path towards explanations about interactions across increasingly autonomous components in industrial IoT. We believe that this will significantly increase explainability in pursuit of GS1-relevant use cases in the domain of industrial automation and technical industries.

Outside of the progress on the research side, we also had to deal with the global challenges due to the Covid-19 pandemic. International research exchanges were put on hold, budgets were reduced, and we could not recruit and retain people from certain geographies because of the travel restrictions. While the pandemic is not over yet, we are grateful that Switzerland has not been badly hit yet and that we found good measures to adapt to the new normal. For example, we successfully held PhD defenses and workshops remotely, are focused on remote collaboration tools to keep in touch, and even could successfully conduct two hackathons in a hybrid & online format.

We keep up the spirit of research & innovation and try to keep up the good collaboration with our colleagues of the global GS1 & Auto-ID Labs community to get through the current times.

Klaus Fuchs & Alexander Ilic

[redacted]

# Project Reports

Using client-based applications for food purchase interventions

# Scaling Nutri-Score Applications

**Jie Lian**
Ph.D. candidate and doctoral researcher

Over the last year, we assessed the potential of automatically generating Nutri-Score values via data science from existing EU1169 data assets and scaling Nutri-Score applications in online shopping environment.

**Motivation for Online Nutri-Score Applictaions**

To promote healthy food choices, front-of-pack (FoPL) like Nutri-Score (NS) have been found effective and became mandated in some countries such as France, but still remain voluntary in most regions of the world [1]. Due to many complicit factors, it can take long to implement mandatory effective FoPL such as Nutri-Score. Besides, with the fact that food producers are not required to publish relevant measures such as fruit-and-vegetable share (FVN) FoPL still require manual categorization and manual labelling by experts for each of the millions of food products that exist. In this case, we are facing to issues: 1) for products with known Nutri-Score, how can they be displayed to the public in existing online shopping envirenment; 2) for products with unknown Nutri-Score, how can we enable the automatic generation of correct Nutri-Score labels from public food composition data. To solve the above problems, we proposed an Nutri-Score extension and an automatic generation algorithm.

**System Design and User Study**

We designed a Chrome brower extension that automatically calculates and displays Nutri-Scores on food product websites when consumers are viewing the websites that present a respective product. A user study was implemented to verify the effects of the Nutri-Score integration on healthy food choices in a simulated grocery eCommerce environment.

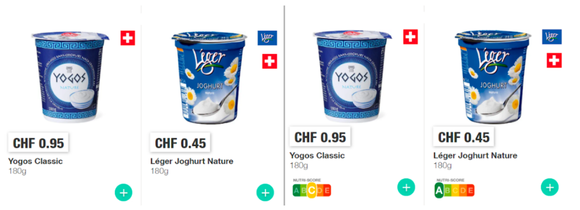


*Figure 2: Left: products without Nutri-Score extension being active; Right: products with Nutri-Score extension being active.*

During the shopping process, the extension enriches the existing webpages from the retailer's eCommerce system with the appropriate Nutri-Score label that is integrated into the respective pages. This is enabled by obtaining products' Global Trade Item Numbers (GTINs) and resolving them to nutritional information through an food composition database. With respect to enabling the usage of this application within user studies, the system furthermore contains mechanisms to create user IDs, assign users to control and treatment groups, track users' shopping process, and to display introductory and post-study information and questionnaires. The collected data from interactions with users is persisted in real-time.

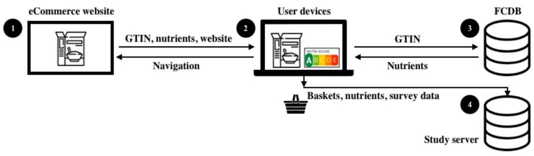


Figure 3:  Architecture and data flow

The user study was conducted in two universities with 135 participants who were randomly into control (CG) and treatment (TG) groups. During the experiment, participants were required to finish an demographic questionnaire then started to do their weekly grocery shopping on Migros website. Once participants submitted their shopping basket, they were redirected to a post-task online questionnaire that assessed specific dietary restrictions of the participants and multiple aspects regarding trust toward the retailer, self-reported food literacy, and the participant's approval of the Nutri-Score. To assess whether labeling food products with their Nutri-Score leads to healthier immediate product choices, we proposed a set of dependent variables including the proportion of healthy food products, the proportion of unhealthy food products, the average Nutri-Score of the shopping basket and the Healthy Trolley Index (HETI) to compare the participants' shopping decisions in terms of nutritional quality.

We assessed the overall impact of the intervention by comparing the selected food items purchased by the control and treatment group respectively. The data analysis shows that on average, the treatment group that was exposed to the digital Nutri-Score food label purchased healthier shopping baskets, as indicated by the significantly higher mean Healthy Trolley Index (HETI) and significantly lower quantities of saturated fat, sugar and unhealthy sugar (sugar quantities in purchased food items except for fruit and vegetables). Meanwhile, there exist further promising, albeit non-significant tendencies that suggest healthier food choices in the treatment group, such as a higher share of healthy products (Nutri-Score A and B) and a lower share of unhealthy products (Nutri-Score D and E). In addition, also dietary fiber, protein and sodium, the products selected by the treatment group feature healthier nutritional compositions in the treatment group. We also explored the impact of digital food labels on the shopping behavior of consumers with low food literacy as they represent an at-risk population for diet-related diseases, along with the effect of conscious perception of the Nutri-Score during the shopping journey. The assessment of participants' perception of the Nutri-Score and of their shopping experience was also included. All results can be found in our paper [3].

**Automatic Generation of Nutri-Score Labels**

In the Nutri-Score generation model, we require the knowledge of product' category and fruit-vegetable-nut (FVN) share since they are critical features to calculate the final value. For the categorization and FVN prediction prediction, we applied three different TF-IDF approaches; and a Na¨ıve Bayes and Random Forrest approach. Then the Nutri-Score calculation algorithm was employed to get the final prediction values. We evaluated the resulting combinations based on precision, recall, accuracy, F1 scores, and their MCC score. Compared to a dummy baseline, we improve accuracy from 0.860 to 0.963 for categories, 0.829 to 0.911 for fruit/vegetables/nuts, and 0.784 to 0.951 for NS scoring with our best model. The study reveal some potentials that by leveraging machine learning to support food composition databases in augmenting already existing product data for correct calculation of the NS will save significant manual effort for product databases, producers retailers in adoption of the Nutri-Score.

**References**


[1] WHO. 2018. Disease prevention by nutrition recommendations.http://www.euro.who.int/en/health-topics/disease-prevention/nutrition. Accessed: 2019-10-11.

[2] Chantal Julia and Serge Hercberg. 2017. Nutri-Score: Evidence of the effectiveness of the French front- of-pack nutrition label. *Ernahrungs Umschau*, 64(12):181–187.

[3] Fuchs, K., Lian, J., Michels, L., Mayer, S., Toniato, E., & Tiefenbeck, V. (2020). Effects of Digital Food Labels on Healthy Food Choices in Online Grocery Shopping. *International Conference for Information Systems (ICIS)*. (submitted)

Unlocking the Power of Customer Data

# Personal Digital Receipt Streams as Enabler for New Service Endpoints

**Klaus Fuchs**
Associate Research Director

### Motivation for Digital Receipts

There exist several advantages of digital receipts over paper-based printed receipts. First of all, the production of paper-based receipts requires trees and therefore prevents significant captures of CO2 in comparison with digital receipts. It was estimated that in Sweden alone, 1.5bn paper receipts can be replaced by digital receipts, leading to savings of 60'000 trees per year [10]. Another study estimated that in the UK, 53'000 trees can be saved by replacing 11.2bn receipts every year [11]. Finally, in the US, savings of 12.4M trees annually are estimated [12]. Handing out paper receipts is also increasing labour costs – a report of the Finish government estimates that 800M Euros could be saved in labour costs if digital receipts could be distributed to consumers automatically [1]. Moreover, paper receipts are usually printed via thermal process contain bisphenol A (BPA), which is a harmful chemical. Besides being un-sustainable and harmful, the most thermal-printed paper receipts fade away quickly, resulting in uselessness for the consumer for warranty purposes or product recalls (Sorensen, 2019). Digital receipts on the other hand, allow consumers to take advantages of warranties and get informed of product recalls independent of a paper sheet with limited lifetime. Finally, digital receipts allow for prevention of tax fraud, today often found in the gastronomy industry, where paper receipts are sometimes not printed, and revenues as well as value-added tax understated.

### Digital Receipt Interest Group

Together with Staffan Olsson from GS1 Sweden, Klaus Fuchs from Auto-ID Labs ETH/HSG initiated the Global Digital Receipt Interest Group (DRIG) to consolidate the discussion on digital receipts across all GS1 MOs. With GS1 colleagues from various countries including Australia, Czech Republic, France, Germany, Switzerland, United Kingdom, United States, the DRIG brings together experts to shape a wholistic perspective on the digital receipt domain. If you are interested to join the discussion, you can find relevant documents in the GS1 community room: https://xchange.gs1.org/cr/sandi/drig/ . The agenda for the group includes a common position paper that details potential areas of interest where GS1 can contribute via the provision of implementation recommendations or potential standards in the future. Such potential contributions can range from the digital receipt data object to interfaces between traceability (e.g. EPCIS) and digital receipts for personalized recalls on mobile phones of end consumers.

### FoodCoach: Funding from Swiss National Science Foundation and Korean National Research Foundation

In December 2019, our research on digital receipts was rewarded with a four-year funding grant within the Korean-Swiss Science and Technology Programme. Our interdisciplinary Swiss-Korean research collaboration with Auto-ID Labs KAIST entitled 'FoodCoach' proposes a novel, scalable and tailored approach towards diet monitoring and interventions. Therefore, we plan to apply artificial intelligence to process customers' automatically collected digital receipts from grocery purchases in order to i) estimate households' and individual dietary behavior, and ii) tailor adaptive interventions to participants, based on their purchase behavior and estimated nutritional context, in order to support healthier food choices.

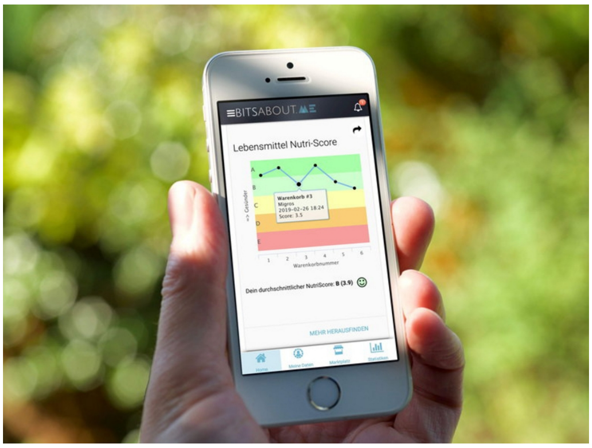

Figure 4: Early-stage screenshot of the FoodCoach functionality that automatically interprets consumers' digital receipts from a nutritional perspective

Thereby, 'FoodCoach' aims to overcome the contemporary limitations of modern diet-related mobile health applications (mHealth). Simultaneously, the proposed project identifies potential adoption drivers (e.g. convenience) and barriers (e.g. privacy concerns) by exploring people's lived experience of digital receipt-based dietary tracking based on survey and focus group research. This affords the opportunity for users' responses (but also non-users' responses) to co-shape the set-up of the study, the design of the interfaces, as well as the kinds of information and hence interventions provided. The project aims to compile the largest digital receipt-based diet panel globally and assess this novel approach's accuracy, scalability, efficacy and ability to reach previously uninvolved users, an important prerequisite in the mitigation of diet-related diseases.

We are currently developing both, the machine learning classification model and the mobile application for tailored interventions in parallel.

We are confident to soon be able to publish promising results from both work streams. To date, N=240 users have participated in the FoodCoach panel.

Figure 4: The FoodCoach mobile app will automatically process digital receipts of a user and tailor recommendations which products are healthier alternatives.

We also received support from Helsana, one of the largest health insurances in Switzerland, to develop the mobile application within the FoodCoach research project. Our goal is to distribute the mobile application to at least 500 users in the next months to learn if such interventions can lead to healthier lifestyles, especially among the at-risk population (i.e. diabetes, overweigh/obese) users. Register at http://www.foodcoa.ch/ to receive updates about the FoodCoach research project.

## References


[1] The digitalisation of receipts to boost the move to real-time economy, Finish Government, Press release, 2019

[2] www.beatthereceipt.com

[3] https://greenamerica.org/skiptheslip

[4] www.verifone.com/en/us/press-release/verifone-introduces-ez-receipt

[5] The eReceipt standardisation application, Finish Government, 2019

[6] http://www.fiscal-solutions.net/index.php/fiscal-countries/

[7] Borderless Real-Time Economy (RTE) Spearhead: eReceipt, EU Interreg Baltic Sea Region, DIGINNO, Republic of Estonia, 2019

[8] M. Bienias et. al., European Digital Guarantee, 2016

[9] https://digitalreceipts.se/swedish-digital-receipt-standard/

[10] I. Sarenmalm, Would you like your receipt? Sustainability perspectives of consumer paper receipts, 2016

[11] https://metro.co.uk/2019/04/02/11-billion-receipts-use-every-year-cant-recycled-forgotten-environmental-problem-9084597/

[12] https://hackernoon.com/digital-receipts-in-retail-b415fbdfde3f

Computer Vision

# Visual Attention Mapping in Physical Stores

**Avik**
Researcher

**Motivation**

In recent years, the retail industry has made major strides towards equipping physical stores with Computer Vision-based systems, enabling real-time inventory monitoring and analytics. For customers, prototypes of checkout-free stores and real-time Augmented Reality-based shopping assistance come to mind. At the core of these services lies the necessity for a system that can accurately locate and recognize (packaged) products in densely packed scenes such as grocery store shelves.

Modern deep learning architectures, namely Convolutional Neural Networks, typically achieve state-of-the-art results in such object recognition tasks. However, the retail setting poses additional major challenges that are yet to be overcome even by the latest research: the number of different product classes to be recognized is typically in the order of several thousands. and the available training data is extremely limited – one studio-quality image per product. Occlusion, deformation and poor lighting of the products visible on store shelves only worsen the situation.

**Approach**

In this work package, we build on our experience in computer vision to contribute to the field by assessing the potential of new approaches to improve the recognition of packaged products on consumer devices (e.g. mixed reality headsets) via computer vision.

As such, we build on a two-stage deep learning-based pipeline addressing these challenges has been implemented. Given as input an image of a store shelf, the first stage detects all the visible products, which are subsequently individually classified in the second stage. To this end, the thesis work begins by considering a basic pipeline based on the combination of two recent state-of-the-art methods for product detection and classification, respectively. With this basic pipeline as inspiration, the core contributions of this thesis are presented, which can be viewed as a proposed collection of steps to substantially improve the architectures of both pipeline stages. These steps range from intuitive ideas like the integration of background samples during training, to more technical ones such as the introduction of attention-based image encoding methods or augmentation of the training data with synthetic samples generated by Generative Adversarial Networks (namely CycleGAN). Additionally, several post-processing techniques such as nearest neighbour re-ranking based on a variety of local image features are proposed, most notably SIFT features. Finally, the effectiveness of these improvement proposals and accuracies of the resulting pipelines are rigorously evaluated on benchmark datasets using a range of metrics. With these contributions, our current resulting pipeline significantly outperforms the previous state-of-the-art, especially in terms of product classification accuracy. A corresponding publication is under way to publish the results.

Figure 5: Mixed reality headset (HoloLens1) in a retail environment

A demo of our detection and recognition pipeline running on a mobile device only can be seen here: [www.youtube.com/watch?v=AdshSvVO2-k](http://www.youtube.com/watch?v=AdshSvVO2-k)

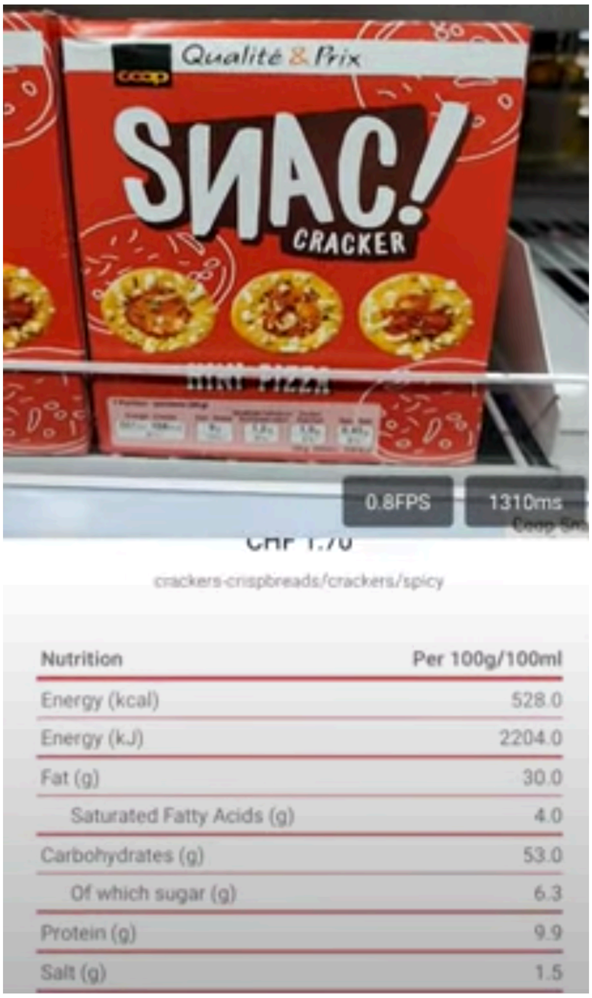


Figure 6: Computer vision based detection can allow barcode-free identification and visualize product characteristics such as nutrients without having to touch a product, potentially especially relevant during the Covid19 pandemic.

With the proliferation of mixed reality headsets and computer vision-based product detection, an increasing number of applications can be developed, from real-time inventory monitoring, customer analytics, to checkout-free stores and real-time Augmented Reality-based shopping assistance.

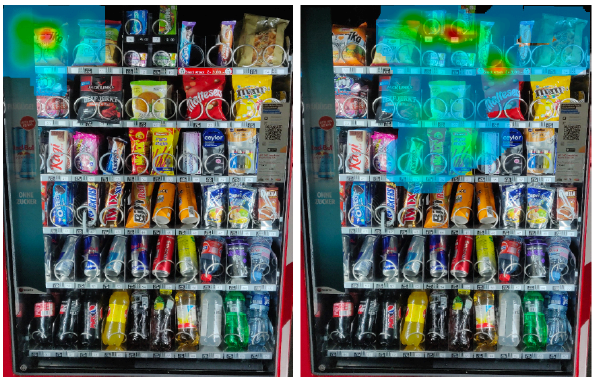

Figure 7: Visual attention mapping via computer vision. Left: Predetermined shopper, right: shopper who is less decisive.

As such, also the real-time detection of user attention becomes feasible. Leveraging tracking of gaze and focal points allows for identification of objects of interest during the purchase process. This could allow for real-time interventions based on the products that a user finds attractive. We will continue to work on and contribute to this upcoming field, which will especially become relevant once Apple, Facebook and Google have released their mixed reality headsets.

## Conclusion and Outlook

Our research exemplifies that in theory, computer vision could substitute barcodes. Still, we do not foresee this happening any time soon, as computer vision-based product detection never yields 100% accuracy, and false positive detections will yield negative consequences (i.e. false prices in self-checkout for example). Still, we expect to see many consumer-relevant applications that will leverage computer vision based product detection and classification as well as attention tracking during customer-product interactions.

Mixed Reality in the Industrial IoT

# Towards Explainable Manufacturing Execution Systems

**Iori Mizutani**
Doctoral researcher

Explainability is one of the emerging topics for Mixed Reality (MR) applications. Fault Detection and Diagnostics tools require explanations about interactions across increasingly autonomous components in industrial IoT settings, such as Manufacturing Execution System (MES). Our research focuses on how we can provide explainability in pursuit of GS1-relevant use cases in the domain of industrial automation, in accordance with the WP3 and the objectives we developed last year.

**Use case: Process Management in Production**

*Figure 1: The PoC MES to produce the Cube Artifacts. (The trajectory of artifacts starts from the left.)*

To identify use cases and assess the scope of explainability in practical environments, we designed and implemented a Proof of Concept MES (Figure 1). The MES produces "Cube Artifacts" – a human worker feeds wooden cubes into the production line. A robot arm equipped with a laser module engraves a QR code on the cube to deliver the artifact. The QR code contains a Digital Link URI for the issued SGTIN, not only for the compartibility with GS1 identification scheme but also to provide cross-domain information retrieval with Web of Things techonology, especially Semantic Web. Figure 2 shows a product page on the EVRYTHNG Dashboard [1] for one such "Cube Artifact" [2].

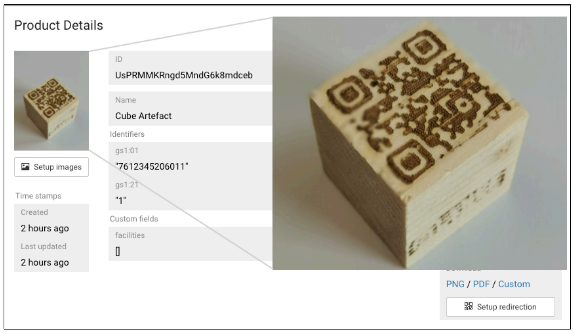


*Figure 2: Each "Cube Artifact" produced in our prototype MES has a laser-engraved QR code of its own Digital Link, which can be dereferenced for available service applications.*

Because the MES is an event-driven system, the relationship between events and outcomes (the production, detection, consumption of, and reaction to events) explains caulsality between components. To model this relationship, we converted the manufacturing process into a dependency graph as depicted in Figure 3 by denoting the entities and the artifacts as finite-state machines. For instance, the production of a single artifact is broken down into 19 steps; the state machine of an artifact transitions between 19 *micro states* (e.g., "the artifact is being conveyed by belt") while *macro states* represent contextual explanations of the MES (e.g., "Laser engraving robot ***L*** is waiting for an artifact ***A*** to be detected by ultrasonic sensor ***U***").

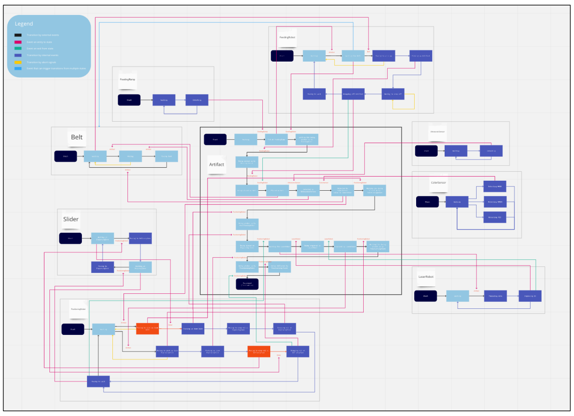

*Figure 3: Dependency graph of the entities in the MES.*

Our knowledge base, inspired by the Semantic Web, stores this dependency graph to be used by other service applications. We created a data model for the graph in OWL, and it contains four major super classes (*Entity*, *Event*, *State*, and *Transition*) and some object properties that interconnect individuals (e.g., "State ***Q*** belongs to Entity ***A***, and it transits to another State ***Q'*** upon Event ***E***"). We also issued GIAIs to entities and assigned them as data properties so that the GIAIs can be used as a key identifier for services based on the GS1 scheme. Our simple Web API [3] enables queries from a MR user interface to elicit latent causal relations between actions and events. Figure 4 is a demo FDD application user interface on HoloLens.

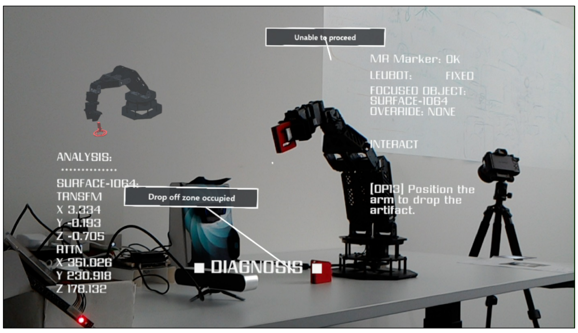


*Figure 4: Our Mixed reality user interface prototype for a FDD application in a robot arm operation.*

## Adaptation of Microservice Architecture

Our underlying motivation for explainable MES is to make complex distributed systems more understandable for human operators by capturing flow dependencies. Industrial automation systems are inherently monolithic due to proprietary protocols or management frameworks. [4] However, this trend has gradually waned in recent years – the information model exchange and the interface between components are becoming more "open" to achieve interoperability among different vendors, as seen for instance in the OPC Unified Architecture standard. This motivates Microservice architectures [5] for MES, having more entities in the system provide abstracted APIs to maintain loose coupling among different services. However, this tracking of interactions across compartmentalized entities becomes more challenging. As seen in Figure 4, our simple PoC system can have tremendously complex causal relations among entities. It is crucial for an explainable system to elicit the necessary information from distributed, separately maintained and deployed entities. Each entity has micro states at runtime, and the combination yields the macro state for the whole system, holding the context of the operation being executed. Reconstruction of the macro states from the runtime would enable useful information overlays to be delivered as a MR application.

## Conclusion and Outlook

Dealing with the distribution to track the state of the system is essentially harder than in monoliths. The internal components of end devices lose visibility because the terminal complexity is hidden by abstraction of the interface and the compartmentalization. We continue to investigate how to resolve this trade-off between modularization and transparency by filling the gap with a semantic graph to provide explanations to the system, leveraging GS1 identification schemes to enable cross-domain information exchange in industrial IoT scenarios.

## References

[1] https://dashboard.evrythng.com
[2] https://id.interactions.ics.unisg.ch/01/7612345206011/21/1
[3] https://api.interactions.ics.unisg.ch/docs/mes-causality
[4] Jammes, François, and Harm Smit. "*Service-oriented paradigms in industrial automation.*" *IEEE Transactions on industrial informatics* 1.1 (2005): 62-70.
[5] Dragoni, Nicola, et al. "*Microservices: yesterday, today, and tomorrow.*" Present and ulterior software engineering. Springer, Cham, 2017. 195-216.

## Explorations beyond core work packages

# GS1 Europe Hackathon: Boosting Honest Products

Besides conducting research within our respective work packages for GS1, the Auto-ID Labs also engage in activities outside the work packages to explore new frontiers: This year, we are supported GS1 Europe's AI hackathon on 'Boosting radically honest products with AI and data' which took place in September 2020. GS1 Europe has reached out to us to support their hackathon, which we gladly did. Concretely, we supported GS1 Europe and Atrify in the preparatory phase of their very first hackathon, and Klaus Fuchs joined the jury board to evaluate the developed prototypes. Together with Google's Cloud Platform, Atrify's (formerly 1WorldSync) product master data APIs, and further challenge sponsors such as Evrythng, we will support the over 177 participants to build revolutionary new solutions on top of product master data.

There were three streams for participants to choose from. First, in the Prototyping stream, participants will develop a real AI-powered solution to address the challenges in the hackathons which are centered around data quality, human nutrition and sustainability. Potential results could include web or mobile applications, websites, a Jupyter Notebook, or a novel API. Second, to cater to participants who might not be able to program, there will be a Concept Stream, in which each group delivers a finished concept that might not include actual technical implementations but be presented in visual form, e.g. as a presentation. Finally, GS1 employees who do not want to participate in the hackathon but want to support the event, can join the Organization Stream.

After the 24h hackathon, a total of 13 applications were submitted. The solutions' focus was centered around healthy nutrition and sustainability. Learn more about the hackathon and please find the submitted solutions under the hackathon's website at:

https://gs1hackathon2020.devpost.com/.

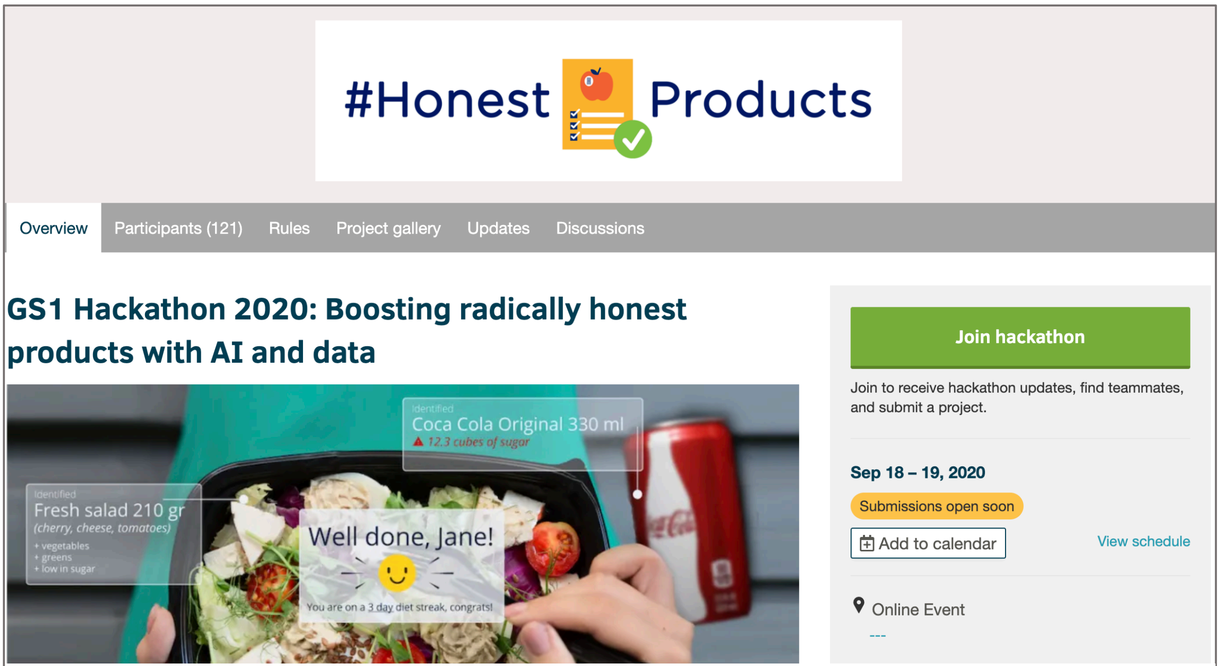

Figure 8: GS1 Europe Hackathon on Product Master Data APIs with Atrify, Auto-ID Labs, Evrything, Google, etc. Learn more at https://gs1hackathon2020.devpost.com/

Explorations beyond core work packages

# HackZurich Hackathon: From Couch Potato to Fitness Hero

In addition to GS1 Europe's first AI hackathon on 'Boosting radically honest products with AI and data', we also hosted a challenge at HackZurich 2020, Europe's largest hackathon. Although this was the seventh time that Auto-ID Labs ETH/HSG hosted a challenge at HackZurich, the level of innovation and technical quality of the implemented solution never fail to surprise us.

In total, over 1'100 developers (300 in person, over 800 online) programmed over 40 hours to develop a total of 145 submitted solutions. The overall winner of Hackzurich2020 was 'AI-Search for spoken audio' which used machine learning to make podcasts searchable. We also were fascinated by 'Clippy', an application that uses historic loyalty card data provided by Migros (largest Swiss retailer and main sponsor of HackZurich2020) to recommend tailored products that are either more sustainable than a user's current habits. Next, 'MR Shopper' used HoloLens2 to support users during food shopping with augmented reality hints. An innovative approach by 'Active4Swag' leveraged the Google Calendar API to implement walks in between meetings during period of nice weather. An app called 'MyClimatePal' rewards users for taking climate-friendly actions, such as choosing public transport or buying sustainable grocery products.

Learn more and find the submitted solutions under the hackathon's website at [https://hackzurich2020.devpost.com/project-gallery/](https://hackzurich2020.devpost.com/project-gallery/)

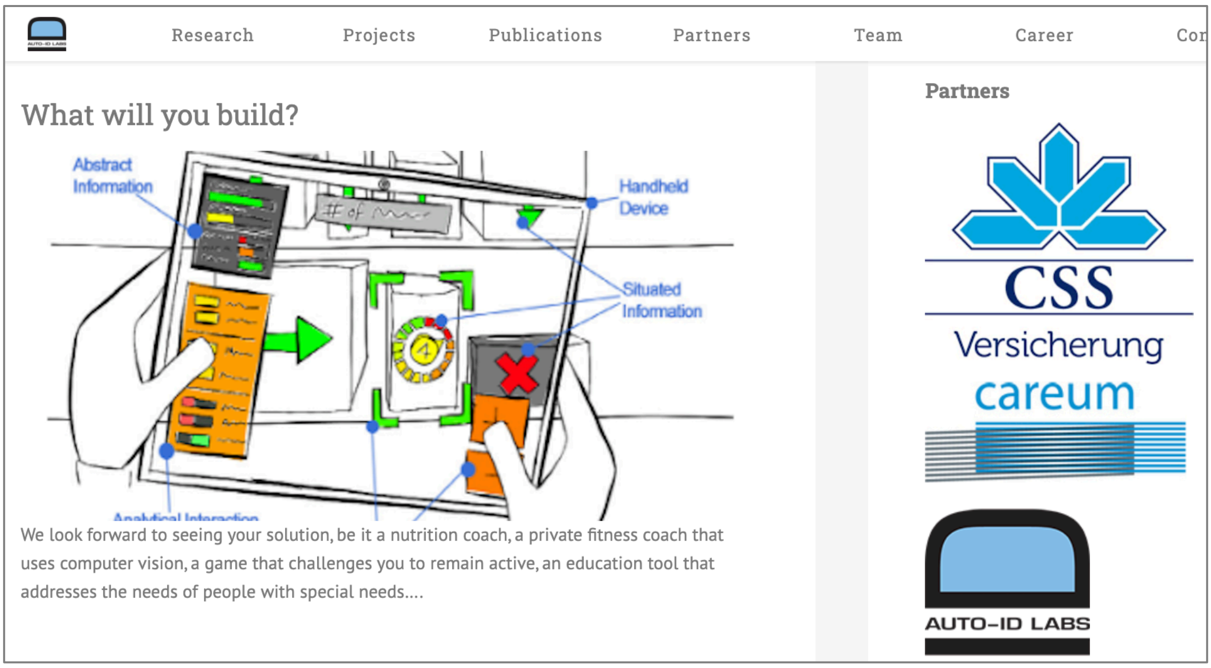


Figure 9: Auto-ID Labs workshop at HackZurich2020 in collaboration with CSS health insurance. Find out more at [www.css.ch/hack](www.css.ch/hack)