

Towards Crowd-Based Requirements Engineering for Digital Farming (CrowdRE4DF)

Eduard C. Groen^{*†}, Kazi Rezoanur Rahman[‡], Nikita Narsinghani[‡], Joerg Doerr^{*‡}

^{*}Fraunhofer IESE, Germany, {eduard.groen, joerg.doerr}@iese.fraunhofer.de

[†]Department of Information and Computing Sciences, Utrecht University, Netherlands

[‡]Department of Computer Science, University of Kaiserslautern-Landau (RPTU), Germany

Abstract—The farming domain has seen a tremendous shift towards digital solutions. However, capturing farmers’ requirements regarding Digital Farming (DF) technology remains a difficult task due to domain-specific challenges. Farmers form a diverse and international crowd of practitioners who use a common pool of agricultural products and services, which means we can consider the possibility of applying Crowd-based Requirements Engineering (CrowdRE) for DF: CrowdRE4DF. We found that online user feedback in this domain is limited, necessitating a way of capturing user feedback from farmers in situ. Our solution, the *Farmers’ Voice* application, uses speech-to-text, Machine Learning (ML), and Web 2.0 technology. A preliminary evaluation with five farmers showed good technology acceptance, and accurate transcription and ML analysis even in noisy farm settings. Our findings help to drive the development of DF technology through in-situ requirements elicitation.

Index Terms—Crowd, crowd-based requirements engineering, digital farming, speech-to-text.

I. INTRODUCTION

The agricultural revolution has brought significant transformations to the domain. The introduction of guidance and sensing systems, telematics, and data management has enabled precision farming, while the green revolution has helped to achieve high productivity gains [1]. More recently, smart farming has emerged, where physical products are supplemented by non-physical services in an agricultural ecosystem, which may include unmanned autonomous robotic and AI-based decision-making systems [2]. Using remote sensing, GPS, weather forecasting, and other technologies, farmers can make more informed decisions, automate agricultural work, make better predictions, and make their business processes more efficient. A collective term for these software-driven innovations is *Digital Farming* (DF) [3]. Although the adoption and acceptance of new technology among farmers remains a limiting factor regarding the innovation potential in DF, development companies in the agricultural sector increasingly have a need to consider feedback from farmers when iteratively improving their (hardware-based) DF equipment and (software-based) DF applications. Farmers can provide contextual information, discuss their ideas, needs, challenges, and requirements, provide feedback on the actual systems they use, provide data for data-driven farming practices, and are a source of other relevant knowledge. As a result, Requirements Engineering (RE) can play an important role in providing solutions for eliciting such feedback from farmers about the technology they use, and in helping development companies improve and innovate their DF products.

Passive data collection methods, typically based on telematics data, are already in use, but these fall short in obtaining context-specific requirements from the farming community. It is likely that many requirements will be missed, which in turn may result in issues ranging from lower adoption rates of new technologies, wasted resources, inefficient solutions, and sustainability failures to complete market failure. Therefore, this knowledge needs to be augmented by feedback directly obtained from farmers. Farming is an international occupation performed by a diverse group of practitioners who use a common pool of products and services, so farmers can be considered a *crowd* [4]. For this reason, it is worthwhile considering whether *Crowd-based RE* (CrowdRE) [4] approaches are suitable to support farmers in communicating their feedback and development companies in using this feedback to further improve their products.

Research questions: We investigate the application of CrowdRE for DF (CrowdRE4DF) by considering its unique domain challenges and designing appropriate solutions through two consecutive studies with associated research questions:

- **RQ1.** What RE-relevant information is contained in online user feedback about DF products?
- **RQ2.** How can user feedback from farmers be successfully captured?

Research contributions: With this work, we make the following contributions:

- 1) We introduce *CrowdRE4DF* as a new subdiscipline of CrowdRE, and characterize it based on early empirical evidence.
- 2) We present a research prototype of the application *Farmers’ Voice* for eliciting speech-based user feedback from farmers, with which we performed an early validation.
- 3) We provide the code for *Farmers’ Voice*, templates, and results from our studies as open artifacts in our online appendix [5].

Paper structure: Section II presents our exploratory study into online user feedback, Section III describes the development and evaluation of our in-situ feedback elicitation tool *Farmers’ Voice*, and Section IV discusses our RQs based on our findings. In Section V, we introduce the domain of DF and discuss relevant work on DF in RE, CrowdRE, and speech-to-text (S2T). Section VI outlines our research plan, and Section VII concludes this paper.

II. STUDY I: EXPLORATORY ANALYSIS OF USER FEEDBACK

The prevailing analysis approach in CrowdRE concerns online user feedback [6]. Thus, our first study investigated the value of online user feedback on DF products. Online user feedback about smart products—such as smart equipment in DF—may differ from other types of applications [7]. User feedback may also be found in farming-specific online channels; this content may differ in terms of focus and language. Obtaining an understanding of these characteristics is helpful in determining how CrowdRE can best be tailored to DF.

A. Methodology

To establish the relevance of user feedback in DF for RE, we identified relevant sources of user reviews (Section II-A1), prepared the data (Section II-A2), and chose appropriate dimensions for classification (Section II-A3).

1) *Selection of Products and Data Sources:* DF products can be distinguished into (a) *DF equipment*, which includes sensors, tractors, and farming bots, as well as intelligent devices (e.g., Yara N-Sensor, Ecorobotix AVO), and (b) *DF applications*, which comprise software that is standalone or used in conjunction with hardware equipment. We determined that online user reviews about DF equipment are rare and typically not available through public channels. For DF applications, in addition to app stores (mainly Google Play), six websites provide publicly available user reviews: [Farms.com](#), [Farmers Weekly](#), [CropLife](#), [Capterra](#), [G2.com](#), and [GetApp](#). We considered 24 farming applications (see Table I). Seven others were omitted, as they had received fewer than ten English-language reviews or their users are not farmers (e.g., crowdfarming applications). Two categories were the most prevalent: (1) *Precision agriculture* applications help to create farm maps, scout the farm, monitor installed sensors, etc., and (2) *farm management* applications provide insights and help to maximize yield and profits.

2) *Data Pre-Processing:* The online user reviews for all selected applications amounted to a dataset of 6,280 reviews, after removing 32 duplicates. Before sampling, we omitted reviews with ≤ 10 characters because these have little informational value (-1,941 reviews) as well as reviews that are written in a foreign language, contain spurious words, or consist predominantly of special characters or emojis (-703 reviews). Because the purpose of this study was to explore typical user feedback about DF applications, we created a bootstrapped random sample of 100 reviews per application to diminish the imbalance in the dataset, which resulted in a sample of 1,335 reviews (see Table I).

3) *Review Classification:* The sample of reviews was manually annotated and reconciled by two authors. Because of the exploratory nature, we began with open coding over a sample of the longest reviews (≥ 500 characters). We determined that the relevant content can be distinguished into the three perspectives listed below. Content that did not fit these three categories was earmarked for discussion, but this did not lead to any new classes. Because a review might address several topics, multiple classes could be assigned. Reviews without clear or relevant content were classified as “None”.

TABLE I
DF APPLICATIONS, ASSOCIATED NUMBER OF UNIQUE REVIEWS, AND NUMBER OF REVIEWS IN THE SAMPLE IN STUDY I.

DF Application	Reviews	Sample	Category
AgriApp	994	100	Farm Management
AgriBus-NAVI	43	30	GPS Guidance
Agrio	108	79	Precision Agriculture
Agrobase	328	100	Precision Agriculture
Agroptima	10	5	Farm Management
Atfarm	14	12	Precision Agriculture
BharatAgri	993	100	Farm Management
BoosterAGRO	13	4	Farm Management
Breeding Wheel	37	23	Livestock Monitoring
Climate FieldView	78	73	Precision Agriculture
CropX	9	7	Precision Agriculture
FarmLogs	148	100	Farm Management
Field Navigator	477	100	GPS Guidance
FieldBee	95	74	GPS Guidance
GoHarvest	61	43	Vendor-Specific
GPS Field Area Measure	1,049	100	GPS Guidance
Landwirt.com Traktor Markt	16	5	Marketplace
OneSoil Scouting	70	51	Precision Agriculture
Plantix	999	100	Precision Agriculture
Soil Sampler	69	37	Precision Agriculture
tractorpool	12	7	Marketplace
xarvio Field Manager	23	18	Precision Agriculture
xarvio Scouting	522	100	Precision Agriculture
Yara CheckIT	112	67	Precision Agriculture
Total	6,280	1,335	

- 1) **System:** A feature, quality, or other technical aspect of the DF application, such as praise, criticism, or requests.
- 2) **Operations:** Use cases regarding the value that the DF application brings to the day-to-day farming business.¹
- 3) **Customer Support:** Experiences receiving assistance or advice from the developer of the DF application, a specialist, or a dedicated community.

B. Results

Of the 1,335 reviews on DF applications, 547 (41.0%) were considered irrelevant (i.e., “none”). Thus, if the reviews with ≤ 10 characters were included, the share of relevant reviews would likely fall below 50%. Figure 1 shows the distribution of the classes. In total, 674 reviews (85.5%) address the *system*, with 583 (74.0%) of them doing so exclusively. These can be analyzed using existing classifiers in RE. A still considerable share of 18.0% (9.0% exclusively) relate to farming *operations*, and 8.6% (5.3% exclusively) mention *customer support*.

User reviews classified as *system* most often request particular crops or livestock types, machines, languages, or geographical areas to be added. For example: “*Major agriculture economy India is not present. What is the reason? Not able to obtain data from the Indian govt?*” Reviews also positively or negatively state whether expectations are met regarding the usefulness, accuracy, and currentness of map data; for example: “*You say that every 3-5 days new satellite [sic] photo comes, but how come I see only 2015 or even older photo instead?*”

¹Works that classified general-purpose applications used classes such as *usage information* or *functional suitability* (e.g., [7], [8]) to annotate descriptions by users of how and in what context they use the application. Due to the specialized nature of farming, we assert that the non-technical and domain-specific descriptions of farming operations are markedly different.

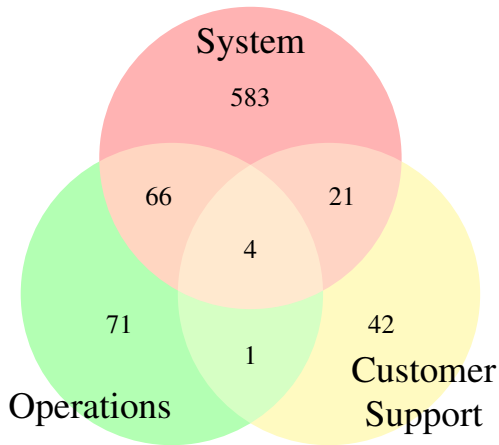


Fig. 1. Distribution of the multi-label classifications.

Idea quite good, but realisation not even close. Usability-related feedback was mostly positive, while other qualities only seemed to be addressed if the application did not start or crashed, failed to exchange data, was not responsive enough, used up too much bandwidth, or had login problems.

User reviews classified as *operations* were typically positive and described the value of the application, which might reveal unique value propositions as well as unidentified potential for deployment in organic farming, for accurately documenting aspects, for taking measurements, etc. For example: *“While it’s true this app is designed more for row crop farms, I use it on our hay ground to keep track of rainfall.”*

User reviews classified as *customer support* mostly praise the timeliness or quality of the responses from the development team, advisors, or online community. There are, however, also negative reviews about the responsiveness or politeness of customer service, such as: *“If you tell them a problem, they will not listen to it for 48 hours. When the plants are completely damaged, their suggestion comes.”* Some reviews also praise the mode of communication (e.g., video calls) or ask how the developer can be contacted.

The strong overlap between *operations* and *system* partly stems from reviews praising the application’s user-friendliness combined with a description of how a farmer uses the application, but more often a farmer motivates a request by explaining their particular situation. For example, a review combining a feature request related to a farm’s location is: *“I am still waiting [sic] offline function. We work near border of republic. There are [sic] no internet.”* Some reviews classified as *system* implicitly reveal a property of the farm, but do not elaborate on the use case behind the request. For a small share of the reviews, we found patterns that indicated that these might be fake, for example when reviews reused sentences from other reviews verbatim, overly promoted the application, or described it as a game. These reviews typically did not contain useful information.

By further analyzing the contents, we found that it is possible to establish a vocabulary of farming-related terms in online

user feedback that could be organized into eight categories: *crops & livestock* (e.g., grain, grow), *property & inventory* (e.g., farm, land), *measurements* (e.g., cost, acre), *weather & irrigation* (e.g., rain, season), *personnel & equipment* (e.g., operator, truck), *GPS guidance* (e.g., field, tracking), *system* (e.g., farm management software, agronomy), and *fertilizers & pesticides* (e.g., chemicals, spray). Only a smaller subset of these terms are directly associated with DF applications or smart equipment, while the majority of the terms relate to aspects typically associated with farming. Only two user reviews that we found mentioned livestock, while many reviews used terms related to crop growing.

C. Threats to Validity

We discuss the main validity threats according to Wohlin et al. [9]. *Construct validity*: We especially found applications related to crop farming, and cannot exclude that we missed relevant sources or applications (e.g., for livestock, forestry, or fishery; [10], [11]). We believe that our methodology helped to assure the integrity of our classifications and note that it served an exploratory purpose and did not seek to create a gold standard. *Internal validity*: Our findings are based on a small amount of data. We strove to be very inclusive about the applications we considered and tried to obtain as many farming-related reviews as possible. We did consider other applications but determined that the reviews did not suit our study. This included an application that received only reviews in German and a GPS navigation application that predominantly received general-purpose user feedback. *External validity*: This work is only intended to characterize user feedback in the farming domain and the results are not very likely to extrapolate to other professional domains.

D. Intermediate Discussion

Through this study, we gained insights into the feedback-giving behavior of farmers. The lack of DF feedback channels, the absence of reviews about DF equipment, and the low number of reviews each DF application receives (see Table I) together suggest that most farmers do not write user reviews. Software applications that users interact with typically receive many more reviews [12]. The number of reviews each application received can still be analyzed manually and does not warrant investing time and effort in an automated analysis pipeline [13], which in turn would forfeit the need to tailor CrowdRE to DF. Similarly, it is not worthwhile investing in a taxonomy on DF as we were only able to curate a total of about 6,000 reviews for this domain, while using the annotated data is prone to over-fitting. Nevertheless, a good proportion of the reviews is requirements-relevant [14]. Like domain-unspecific feedback, the feedback highlights functional and quality aspects of the system [7], while valuable insights can also be gained from descriptions of the usage (i.e., farming operations) and the customer support received [8]. In Section IV, we will discuss our key findings of Study I to answer RQ1.

III. STUDY II: EVALUATION OF A SPEECH-TO-TEXT FEEDBACK SOLUTION

Study I revealed that due to the scant amount of user feedback farmers generate, automated user feedback analysis methods will not contribute much to a better understanding of this crowd of users. Discussions with domain experts allowed us to conclude that this is caused by the farmers' work being demanding. They lack the time to stand around typing a review on their smartphone and do not have much opportunity to type out a user review on a computer because they are not often at their desk. This is unlikely to change through motivational measures encouraging farmers to write more user reviews.

This finding gave rise to Study II, where we set out to design and evaluate a solution that takes into consideration the practical limitations faced by farmers while exploring other suitable ways of capturing user feedback from farmers. Considering their daily schedule, we dismissed any solutions that would require farmers to take the time to provide written feedback. When they are in the office, they are busy preparing or documenting the field work, and out in the field they cannot be reasonably expected to pause for a moment. However, just like they make phone calls during the day, for example while driving their tractor, they could record their feedback through speech in situ. This resulted in our research prototype *Farmers' Voice*; a speech-to-text (S2T) and audio sentiment analysis (ASA) application that allows crowds of farmers to speak with one voice. We evaluated our implementation regarding its viability to help with feedback collection in terms of whether farmers would accept such a solution and whether it would work in a farm setting. The application and all resources necessary to replicate this study are available in our online appendix [5].

A. Analysis & Conceptual Solution

We found that existing S2T applications and smart voice assistants do not support integration into a pipeline, but determined that common S2T features include multi-language support, live transcription, recording control buttons (i.e., start, pause, stop, reset), offline mode, data privacy measures, user-friendliness, transcript editing, and tutorials.

To counteract the limited literature on DF and extend the research team's domain knowledge, we consulted with an agriculture economist who has been researching DF for four years and seasonally works on the family farm. He recommended postponing liaising with farmers until we could provide them with tangible results because they are often critical before adopting new technologies. His most important recommendations are: (1) It is crucial to farmers to be able to contact support personnel, e.g., technical contacts at their equipment dealer or agricultural cooperative and dedicated account managers at companies developing their business software. Farmers mostly call them, which further supports our choice of a speech-based solution. (2) The likelihood that farmers will be able to provide user feedback while sitting on a tractor will only increase once automated steering systems become more common in modern tractors. Any solution should be able to handle the noise levels and work without an Internet

connection. (3) The use of smart voice assistants may be on the rise, but farmers often lack the time to interact with them. They are more likely to be willing to send a one-off voice message, especially when short feedback cycles show a benefit of investing their time.

We considered three types of *stakeholders* for *Farmers' Voice*: farmers, support personnel, and requirements engineers. In this paper, we focus on the *farmers*, who will use the application to record and submit feedback about DF applications and DF equipment, such as (technical) problems, expectations, or ideas. Their willingness to use the application is central to its success. The feedback should then be sent to the *support personnel*, e.g., technical contacts or account managers, who can respond directly in case of a technical problem and prioritize the feedback. The *requirements engineers* can then analyze the user feedback for recurring patterns, needs, and ideas.

Our analysis of S2T applications, our consultation with a domain expert, and iterative discussions among the authors resulted in eleven requirements for our research prototype, which are listed in Table II. These were translated into use cases, which were then used for our interface design in *Justinmind*. The S2T and ASA modules allow us to determine whether farmers prefer reviewing and submitting a text-based transcript or an audio recording augmented by sentiment and keyword indicators. The baseline feedback module was designed to store data differently and provide additional reporting functionalities necessary for evaluating the application's performance against predefined benchmarks. The user journey in the application as shown in Figure 2 is as follows: After logging in (screen 1), the user encounters a landing page where recordings can be made and managed in free-form feedback mode and S2T processing (2). The transcript is shown directly on-screen (3). In baseline mode, a report can be accessed (4&5). The user can activate the ASA module in the tabs of the top navigation bar (6). The ASA module includes a media player for each recording and a menu to display the detected emotions (7), and the ability to upload recordings (8).

B. Software Architecture & Implementation

We developed *Farmers' Voice* in *Visual Studio Code* following the *Prettier* conventions, using *Git*, *Github*, and *HuggingFace* for version control, and *Github Actions* for continuous integration and deployment. A high-level overview of the architecture and its components is shown in Figure 3. Users interact with the front-end through a web-based application based on the component-based architecture of *ReactJS*. We integrated Mozilla's Web Speech API using the *react hook* "react-speech-recognition", and reusable user interface components through the *Material UI* library. We designed *Farmers' Voice* as a progressive web application to ensure responsiveness, cross-platform adaptability (e.g., to different screen or window dimensions), and offline functionality to ensure that the web-based application will remain accessible in poor network conditions (e.g., out in the field). Sending the data to the model for processing over APIs does require an Internet connection. Switching languages

TABLE II
REQUIREMENTS FOR THE RESEARCH PROTOTYPE OF *Farmers’ Voice*.

ID	Requirement
RQ01	The application shall have a user interface that non-tech savvy users are able to operate.
RQ02	The application shall allow the user to record feedback through speech and see the generated text-based transcript.
RQ03	The application shall allow the user to submit their feedback in the form of raw audio clips or text-based transcripts.
RQ04	The application shall be cross-platform to enable its use on as many platforms as possible, with an emphasis on mobile platforms.
RQ05	The application must assure data privacy and include basic security features such as a login system and data removal policies.
RQ06	The application shall process the transcribed text, at least regarding keywords, summaries, and sentiments.
RQ07	The application shall allow the user to switch between German and English for both the interface and speech processing.
RQ08	The application shall allow the user to download the transcript, processing results, and metrics as a report in an appropriate file format.
RQ09	The application shall have an audio sentiment analysis module that processes audio recordings.
RQ10	The application shall allow the user to submit both new and older recordings for audio sentiment analysis.
RQ11	The application shall allow the user to toggle between the standard free-form feedback and a research-oriented baseline feedback module.
RQ12	The application shall be deployed in an online hosting environment to make core components available in remote setups.

and localization was achieved through the JavaScript-based I18n internationalization framework. The front-end stack was deployed in Google Firebase. The back-end provides data to the front-end components and handles post-processing. For natural-language processing, we use HuggingFace’s pre-trained models, which are accessed using Python with Gradio’s “use via API” endpoints deployed in HuggingFace Spaces, through which we invoke Superb/wav2vec2-base-superb-er and distilbert-base-uncased-finetuned-sst-2-english for sentiment analysis, transformer3/keywordextractor for keyword extraction, and knkarthick/meeting_summary for summarization. The DeepL API deployed in Heroku translated transcripts into English. We decided not to employ noise reduction algorithms in the research prototype because their consistency cannot be assured across devices; we tested in the evaluation whether there would be a need to implement such functionality during live transcription or when processing audio files.

C. Evaluation

We evaluated *Farmers’ Voice* to measure its performance in realistic settings and determine the acceptance of our solution among prospective end users. We needed to make sure that the application will work in farm settings and would be used by farmers at all, before our research addresses the quality of its data processing of the user feedback content.

1) *Participants & Instrumentation*: The evaluation was conducted at Hofgut Neumühle, an associated farm specializing in research and training for livestock farming. Through its network, we solicited five farmers (one female, four male) from the German state of Rhineland-Palatinate who were available to participate on 2 August 2023: three experienced grain farmers (P1, P2, and P5), a novice vegetable farmer (P4), and an intermediate pasture farmer (P3). Regarding their technological proficiency, they were all very familiar with mobile applications and would feel comfortable using a speech-to-text application. Except for P4, they actively used technology professionally and had prior experience with voice assistants like Siri or Alexa. Developers of DF technology typically target farmers with these properties to market their solutions to. Most participants considered themselves rather good typists, except P3, who is a slow typist. However, P2, P3, and P5 stated that they would prefer giving feedback using speech, while P1 and P4 would

additionally like to provide written feedback. The experiment was conducted using a OnePlus 7 Pro smartphone with Android 12 installed, which was tested prior to the experiment to prevent technical issues such as connectivity problems.

2) *Procedure*: The evaluation was conducted in three stages: (1) briefing & pre-evaluation survey (30 minutes); (2) field experiment (75 minutes), and (3) focus group & debriefing (100 minutes). In the briefing, we introduced the participants to the research and familiarized them with the *Farmers’ Voice* application, its modules and workflow, and clarified questions. In the pre-evaluation survey, which is included in our online appendix [5], we elicited their demographics and expectations.

We performed a 2 (noise conditions) \times 2 (feedback types) within-subjects field experiment. The same steps were performed in all scenarios by all five participants to ensure that any differences could be attributed to the change in the environment. Accordingly, the experiment consisted of 20 treatments. Two noise conditions—office setting and field setting—simulated realistic scenarios in a farmer’s day. The *office setting* was located in a room at the farm and resembled a low-noise environment. For the *field setting*, the farmer was seated on a tractor running stationary but with the power set to maximum (i.e., generating the loudest possible noise), to represent a high-noise environment. This allowed us to compare transcription accuracy. Two feedback types—baseline and free-form—determined the content. In the *baseline feedback* condition, each participant was provided with the prepared baseline text included in our online appendix [5] to enable quantitative comparisons. They read this text aloud while the application was recording. In the *free-form feedback* condition, the participants picked a topic, such as a new farming technique, a crop management strategy, or an experience with DF equipment, and freely discussed notable functionalities, ease of use, impact on farming operations, (dis)advantages, and suggestions for improvement. In each condition, a recording was made separately using each module (S2T and ASA). This way, the farmers could get a feel for what it is like to use the application. Their acceptance of the proposed technology and perception of the application’s usability were captured using a focus group to get feedback about the use of the application, satisfaction, perceived usefulness, challenges, and preferences.

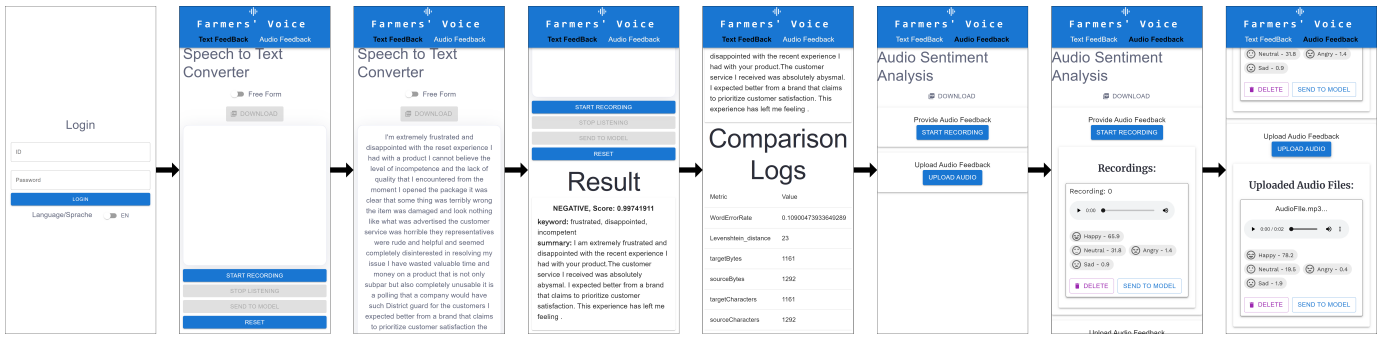


Fig. 2. Screenshots from the mobile layout of *Farmers' Voice* showing a possible user journey.

3) *Data Analysis*: For data collection, *Farmers' Voice* generated a report with several typical S2T metrics. The *Word Error Rate (WER)* [15] measures the overall accuracy of the S2T model. It compares the number of errors (substitutions, deletions, and insertions) in a transcription to the total number of words against the baseline text. The *Levenshtein Distance* or *edit distance* [16] quantifies the differences between the original spoken text and the transcribed text to determine the accuracy of the transcription compared to the original speech recording. It measures the difference between two strings in terms of the minimum number of single-character edits required to transform one string into the other. We also compared the length of each transcription (*target bytes and characters*) to the baseline text—1,597 source bytes and 1,572 source characters—to determine the completeness of the transcription. We calculated the significance of the effect of noise conditions using a paired T-test, which was one-tailed because the field setting was hypothesized to perform worse than the office setting. Because five item pairs yield a low effect size, we set the significance level to $\alpha = 0.10$.

4) *Results*: The performance of *Farmers' Voice* in both noise conditions is shown in Table III. In the field setting, WER and Levenshtein distance are higher. Significance was not reached, but the trend suggests that a difference could be observed with more participants. The lower number of target bytes and characters in the field setting also reflects slightly greater difficulty in transcribing some participants' speech; e.g., with P2, the number of misinterpretations quadrupled in the field setting. However, the performance was not significantly worse in the field setting. This shows that the application is only slightly less accurate in a high-noise setting.

The participants' acceptance of the application increased during the experiment. They were critical in the pre-evaluation survey, expecting an S2T application to capture their feedback accurately and reliably if they were to accept it. In the focus group, all were positive about the use of the application,



Fig. 3. Overview of the architecture for *Farmers' Voice*.

TABLE III
MEANS AND STANDARD DEVIATIONS (SD) FOR THE SETTINGS AND PARTICIPANTS OF THE MEASURED S2T PERFORMANCE IN THE BASELINE MODULE, AND THE SIGNIFICANCE OF THE DIFFERENCE (p); THE OMITTED VALUES FOR DIFFERENCE EQUAL THE TARGET VALUES.

Metric	Office Setting		Field Setting		Comparison		
	Mean	SD	Mean	SD	Mean	SD	p
Word Error Rate	0.30	0.07	0.36	0.08	0.05	0.01	0.1109
Levenshtein Distance	67.2	15.1	79.4	17.2	12.20	2.12	0.1040
Target Bytes	1422.8	105.7	1350.6	158.9	-72.20	53.20	0.2180
Byte Difference	-174.2	105.7	-246.4				
Target Characters	1398.8	104.9	1329.0	157.8	-69.80	52.88	0.2242
Character Difference	-173.2	104.9	-243.0				

with expectations exceeded for some. They found that the transcription worked intuitively and fast, and returned accurate feedback, with P1 saying: “During free-form feedback, I said many different words and I was impressed with the performance.” They did notice that there were problems with common farming abbreviations like NIR and NDVI,² which should be supported to handle reports of complex topics. They also expressed concerns about the transcription handling heavy dialects, of which there are almost 250 in Germany alone. They recommended addressing both issues sooner rather than later.

The participants agreed that the application was comfortable to work with and that they could provide feedback with ease. This was even the case for P4, who had never used voice assistants. They would find it easier to provide feedback in this way than by contacting their dealer or other support personnel, but P5 noted that emotions can be better expressed through personal communication. All participants indicated that they preferred providing feedback verbally rather than typing it. As a slow typist, P3 found this to be both more convenient and a time-saver. This was in line with the expectations expressed in the survey. This evidence suggests that S2T can be an easier way for farmers to provide feedback. Four participants preferred the S2T module, while P3 preferred the audio sentiment analysis module because “one can draw better conclusions about the customer's mood.” The participants found that the ASA module was not very accurate in identifying emotions other than anger; it tended to label most feedback as either neutral or sad.

In the pre-evaluation survey, data privacy was a major concern. Trust increased during the experiment, but the

²Near-Infrared; Normalized Difference Vegetation Index.

participants agreed that many farmers may be concerned about how the data is used and who will receive it. They provided a compelling example: Farmers might fear that if feedback about their fertilizer use reaches federal or state authorities, they could use it to impose further regulations or restrictions.

D. Threats to Validity

We discuss the main validity threats according to Wohlin et al. [9]. *Construct validity*: We strove to evaluate the most crucial features of *Farmers' Voice*. It was not possible to balance the small number of participants across treatments or draw statistically significant conclusions about the conditions. To diminish instrumentation effects, we used a single smartphone without noise cancellation features. Android limits the microphone listening time of an app, which we circumvented by instructing the participants to avoid longer pauses. *Internal validity*: To mitigate the Hawthorne effect during the experiment or the Rosenthal effect during the focus group due to the development and evaluation being performed by the same researcher, we included a neutral second experimenter and carefully prepared the guidelines, among other things to facilitate an atmosphere in which the participants felt comfortable to speak openly. *External validity*: We evaluated the application with professional farmers; a stakeholder group that is difficult to involve in experiments due to their limited availability. Despite their diversity in terms of gender, experience, technological affinity, and type of cultivation, we cannot exclude a selection bias towards farmers who (a) are from the same region in Germany and (b) might be more accepting of technological innovations. Regarding the latter, the participants were skeptical before the evaluation and were asked in the focus group to extrapolate whether they thought that other farmers would accept the application. Also, a critical attitude towards technology does not mean complete refusal to use a technology. The outcomes of our evaluation may be limited to the domain of farming.

IV. GENERAL DISCUSSION

In this work, we set out to investigate whether CrowdRE can be applied to DF despite aggravating domain characteristics, and what tailoring may be necessary. The limited amount of research on DF required us to address many questions and assumptions along the way. Several challenges may inhibit the success of CrowdRE. Even more than in other professional fields, farmers have little opportunity to give feedback. Especially outside conditions for providing feedback might be unfavorable due to workload, weather, noise, Internet connectivity, etc. Many farmers are reluctant to adopt DF technologies, often out of insecurity over their technical literacy, English proficiency, or the trustworthiness of the application's privacy [17]. Adoption is greater among farmers with larger farms, better support infrastructures, and exposure to recommendations from peers. Privacy concerns and mistrust of governing bodies might also cause farmers to fear that their thoughts and ideas are misused. The user feedback from

farmers might also be harder to analyze automatically due to different languages and dialects, jargon, and abbreviations.

As to whether online user feedback from the farming domain contains RE-relevant information (**RQ1**), we found too little feedback about DF equipment, some feedback on (software-based) DF products especially on Google Play, and few domain-specific or professional channels. RE-relevant information in the reviews is usually about the system, similar to online user feedback about general-purpose software applications, so it can be classified into features, qualities, or other RE-related dimensions by existing means (e.g., [18], [19]). Descriptions of how a DF application is used in farming operations may provide relevant insights to RE (cf. [8]). Although we can positively answer RQ1 based on the relevance of the content of online user feedback, DF applications are currently receiving too little feedback to allow for automated classification. Giving feedback might conflict with a farmer's daily schedule, which might be a unique challenge for applying CrowdRE in DF.

Investigating the question of how to capture user feedback from farmers in the best way (**RQ2**), we adapted our strategy to their daily routine using *Farmers' Voice*, an in-situ solution that allows spoken messages to be recorded, optionally transcribed, and submitted. It performed surprisingly well in a noisy tractor cabin even without the use of noise reduction technology. This means that giving feedback in a time and manner that is convenient for farmers, i.e., hands-free while driving a tractor, is a viable option. The farmers in our evaluation were open to using *Farmers' Voice*. The experiment reduced their initial hesitation and met their expectations in terms of accuracy and usability. They did, in fact, prefer giving verbal feedback rather than written feedback and favored the S2T module over the ASA module. Accordingly, these technologies have the potential to successfully capture user feedback from farmers, thus to realize CrowdRE4DF.

V. RELATED WORK

Digital Farming (DF) refers to connected, knowledge-based agricultural production systems that utilize intelligent network and data management tools to automate sustainable processes in agriculture [20]. Its key components include sensors, Internet of Things devices, data analytics, autonomous robots and drones, and software applications. Value is created by connecting smart machines, with a huge emphasis on interoperability [21]. Farms are often challenged by the costs of implementing new technologies, weak and unstable Internet signals, or low bandwidth on the field. Older generations of farmers often depend on their digital native children for solving technical issues. Ofori et al. [17] found that, accordingly, farmers are quicker to adopt embodied-knowledge technologies (i.e., knowledge that is gained through direct experiences and interactions) than information-intensive technologies. Pfeiffer et al. [22] established that a positive attitude of German citizens towards farming and trust in farmers correlated with a positive attitude towards the perceived benefits of DF technologies, which include improved quality of life for farmers through

reduced workloads, more environmentally-friendly production, and improved animal welfare and overall health.

In RE, DF has received fairly limited attention. Blasch et al. [23] found that the adoption of Precision Farming technologies in Central Italy was moderated by initial investment costs, farm size, and economies of scale. Factors positively influencing adoption were financial support and networking and knowledge sharing among farmers, which demonstrates the important role of recommendations of new technologies from peers. Mannari et al. [24] found that class, goal, and process models were helpful for communication during requirements elicitation in a smart irrigation pilot study in Italy, even with non-technical stakeholders such as agronomists and farmers. Ferrari et al. [11] established that farmers in remote mountain areas benefit from digital technology in terms of efficiency, reduced work, and strengthened social bonds, but that it can also form a distraction, limit problem-solving capability, and impair working conditions.

CrowdRE is an umbrella term for (semi-)automated approaches for collecting and analyzing information from a crowd to derive validated user requirements [4]. It presumes a crowd to be a sizable group of current or potential users with a common interest in a (software) product. Collecting text-based feedback has received more attention than monitoring usage data [25] and mainly involves user feedback analysis [12], [6]. Several mobile applications in RE elicit in-situ user feedback in the form of text and images, notably *MyExperience* [26], *iRequire* [27], *ConTexter* [28], *AppEcho* [29], and *Rich Parking* [30]. Johnson et al. [31] performed the only CrowdRE study to date in the domain of farming by identifying CrowdRE challenges for software ecosystems for the case of *Figured*, a farming financial management system that is a partner application in the accounting ecosystem *Xero*. Farming development companies have an information need, but it is challenging to motivate farmers to provide user feedback.

We found only few demonstrations of S2T applied in RE. Soares et al. [32] proposed the *VoiceToModel* framework to improve the accessibility of the modeling process. With it, people with physical or visual impairments could create simpler goal, conceptual, and feature models by using voice commands. Their architecture demonstrates that such an application requires multiple external components to be integrated. The pipeline by Maghilnan and Rajesh [33] demonstrates that speech-based sentiment analysis additionally requires supporting speech recognition and speaker recognition to distinguish between two or more speakers. For space reasons, we do not discuss the state-of-the-art approaches in S2T research here.

VI. RESEARCH PLAN

Our speech-based feedback solution respects the farmers' daily schedules and has the potential to be accepted among farmers. The participants in our evaluation were satisfied with the processing by the application. A next step is to collect actual feedback from farmers. Together with the stakeholder category of support personnel, we will seek to determine the quality of the feedback received. An open research question is

whether farmers will prefer giving unstructured feedback or being interviewed by a conversational agent asking guiding questions such as (a) "What do you think of the product?" (i.e., *system*), (b) "How do you use the product?" (i.e., *operations*), and (c) "How can we help you better?" (i.e., *customer support*). We also plan to improve the interaction, privacy, and transcript editing functionality of *Farmers' Voice*.

Giving a large crowd of farmers a voice may have a profound impact on their ability to steer decision-making processes of which they otherwise are only at the receiving end. One example are measures to make regions more resilient to the effects of climate change. We are preparing a research project in an area with a high geographic density of farmers that aims to not only passively collect requirements, but to specifically probe for their expertise in the co-creation of solutions. Such an influx of user feedback from farmers will provide us with the opportunity to design a robust CrowdRE4DF analysis pipeline that can achieve an accurate automated derivation of requirements from the user feedback: the central notion of CrowdRE [4]. Realizing such a pipeline is challenging because of the need to augment it with farming-related data that helps support the domain language and farmers' natural dialects. The vocabulary mentioned in Study I is a first step in this direction. Like other applications suggested [27], [28], the pipeline may also need to support multi-modal feedback.

Offering speech-based feedback technology across systems while assuring cybersecurity reveals two limitations that can only be overcome if this is incorporated into domain standards such as the interface of the agricultural communication protocol ISOBUS [34] or within the *GAIA-X* ecosystem. Farmers could then provide feedback directly through such a communication system. This would help assure data sovereignty and processing within a trustworthy environment. With sufficient DF technology such as sensors and Internet of Things devices in place, a context-aware system could also encourage farmers to provide feedback when they detect events such as malfunctions.

Although we will initially cater to farmers' needs, we recognize the potential for pivoting to domains and contexts with similar challenges. S2T technology could allow other stakeholders to give feedback in situ, such as factory workers operating machinery in noisy environments, drivers who should keep their hands on the steering wheel, or underground mine workers dealing with an unreliable communication network. Giving hands-free user feedback supports better inclusiveness of persons with an impaired ability to operate a device due to pain or physical disabilities, or who have trouble reading and writing due to visual impairments or illiteracy.

VII. CONCLUSION

In this paper, we analyzed the potential of applying CrowdRE in DF. There are unique challenges, which, among other things, cause only little online user feedback to be available. Enabling farmers to record feedback during their daily activities shows promise to help drive the development of DF technology. We have shown that farmers are ready to accept the idea and that our application provides accurate results even in noisy

farming environments. This gives DF technology developers direct access to farmers' feedback, while farmers can become the main drivers of technological advances in DF and improve their ability to produce food more sustainably.

ACKNOWLEDGMENTS

The authors would like to thank Daniel Eberz-Eder of the Dienstleistungszentrum Ländlicher Raum Rheinhessen-Nahe-Hunsrück for arranging our contacts, Dr. Christian Koch of the Lehr- und Versuchsanstalt für Viehhaltung, Hofgut Neumühle for providing his facilities, the farmers who generously donated their time to participate in our evaluation, Christoph Merscher for assisting with the experiment, and Sonnhild Namingha for proofreading this paper.

REFERENCES

- [1] E. Lejon and P. Frankelius, "Sweden innovation power: Agritech 2015," 2015.
- [2] J. Dörr, "Digital farming in the context of digital ecosystems," *ATZ Heavy Duty Worldwide*, vol. 14, p. 68, 2021.
- [3] J. Dörr and M. Nachtmann, *Handbook Digital Farming: Digital transformation for sustainable agriculture*. Springer, 2022.
- [4] E. C. Groen, N. Seyff, R. Ali, F. Dalpiaz, J. Doerr, E. Guzmán, M. Hosseini, J. Marco, M. Oriol, A. Perini, and M. Stade, "The crowd in requirements engineering: The landscape and challenges," *IEEE software*, vol. 34, no. 2, pp. 44–52, 2017.
- [5] E. C. Groen, K. R. Rahman, N. Narsinghani, and J. Doerr, "Towards crowd-based requirements engineering for digital farming (CrowdRE4DF): Online appendix," 2024, Zenodo. [Online]. Available: <https://doi.org/10.5281/zenodo.10972870>
- [6] C. Wang, M. Daneva, M. van Sinderen, and P. Liang, "A systematic mapping study on crowdsourced requirements engineering using user feedback," *Journal of Software: Evolution and Process*, vol. 31, no. 10, p. e2199, 2019.
- [7] E. C. Groen, S. Kocpzyńska, M. P. Hauer, T. D. Krafft, and J. Doerr, "Users—The hidden software product quality experts? A study on how app users report quality aspects in online reviews," in *Proceedings of the 25th International Requirements Engineering Conference (RE)*. IEEE, 2017, pp. 80–89.
- [8] M. Anders, B. Paech, and L. Bockstaller, "Exploring the automatic classification of usage information in feedback," in *Proceedings of Requirements Engineering: Foundation for Software Quality (REFSQ)*, D. Mendez and A. Moreira, Eds. Springer, 2024, pp. 267–283.
- [9] C. Wohlin, P. Runeson, M. Höst, M. C. Ohlsson, B. Regnell, and A. Wesslén, *Empirical strategies*. Boston, MA: Springer US, 2000, pp. 7–24.
- [10] D. van der Linden, "Interspecies information systems," *Requirements Engineering*, vol. 26, no. 4, pp. 535–556, 2021.
- [11] A. Ferrari, F. Lepore, L. Ortolani, and G. Brunori, "Eliciting the double-edged impact of digitalisation: a case study in rural areas," in *Proceedings of the 31st International Requirements Engineering Conference (RE)*, 2023, pp. 157–168.
- [12] D. Pagano and W. Maalej, "User feedback in the appstore: An empirical study," in *Proceedings of the 21st International Requirements Engineering Conference (RE)*. IEEE, 2013, pp. 125–134.
- [13] E. C. Groen, J. Schowalter, S. Kocpzyńska, S. Polst, and S. Alvani, "Is there really a need for using NLP to elicit requirements? A benchmarking study to assess scalability of manual analysis," in *REFSQ Joint Proceedings of the Co-Located Events*, K. Schmid and P. Spoletini, Eds. CEUR Workshop Proceedings 2075, 2018, p. Article 11.
- [14] M. Sihag, Z. S. Li, A. Dash, N. N. Arony, K. Devathanan, N. Ernst, A. B. Abu, and D. Damian, "A data-driven approach for finding requirements relevant feedback from TikTok and YouTube," in *Proceedings of the 31st International Requirements Engineering Conference (RE)*. IEEE, 2023, pp. 111–122.
- [15] D. Klakow and J. Peters, "Testing the correlation of word error rate and perplexity," *Speech Communication*, vol. 38, no. 1–2, pp. 19–28, 2002.
- [16] V. I. Levenshtein, "Binary codes capable of correcting deletions, insertions and reversals," *Soviet Physics Doklady*, vol. 10, no. 8, pp. 707–710, 1966.
- [17] E. Ofori, T. Griffin, and E. Yeager, "Duration analyses of precision agriculture technology adoption: What's influencing farmers' time-to-adoption decisions?" *Agricultural Finance Review*, vol. 80, no. 5, pp. 647–664, 2020.
- [18] Z. Kurtanović and W. Maalej, "Automatically classifying functional and non-functional requirements using supervised machine learning," in *Proceedings of the 25th International Requirements Engineering Conference (RE)*. IEEE, 2017, pp. 490–495.
- [19] W. Maalej and H. Nabil, "Bug report, feature request, or simply praise? On automatically classifying app reviews," in *Proceedings of the 23rd International Requirements Engineering Conference (RE)*. IEEE, 2015, pp. 116–125.
- [20] CEMA aisbl – European Agricultural Machinery, "Digital farming: What does it really mean?" 2017. [Online]. Available: https://www.cema-agri.org/images/publications/position-papers/CEMA_Digital_Farming_-_Agriculture_4.0_13_02_2017_0.pdf
- [21] J. Dörr, M.-A. Favier, and F. Möhrle, "Interoperabilität im digital farming: Status quo und zukünftige herausforderungen," *Wirtschaftsinformatik & Management*, vol. 16, 2023.
- [22] J. Pfeiffer, A. Gabriel, and M. Gandorfer, "Understanding the public attitudinal acceptance of digital farming technologies: A nationwide survey in Germany," *Agriculture and Human Values*, vol. 38, no. 1, pp. 107–128, 2021.
- [23] J. Blasch, B. van der Kroon, P. van Beukering, R. Munster, S. Fabiani, P. Nino, and S. Vanino, "Farmer preferences for adopting precision farming technologies: A case study from Italy," *European Review of Agricultural Economics*, vol. 49, no. 1, pp. 33–81, 2022.
- [24] C. Mannari, G. O. Spagnolo, M. Bacco, and A. Malizia, "Digitalisation of agriculture: Development and evaluation of a model-based requirements engineering process," in *REFSQ Joint Proceedings of the Co-Located Events*, A. Ferrari and B. Penzenstadler, Eds. CEUR Workshop Proceedings 3378, 2023, p. PT Paper 2.
- [25] J. A. Khan, L. Liu, L. Wen, and R. Ali, "Crowd intelligence in requirements engineering: Current status and future directions," in *Proceedings of Requirements Engineering – Foundation for Software Quality (REFSQ)*. LNCS 11412, 2019, pp. 245–261.
- [26] J. Froehlich, M. Y. Chen, S. Consolvo, B. Harrison, and J. A. Landay, "MyExperience: A system for in situ tracing and capturing of user feedback on mobile phones," in *Proceedings of the 5th International Conference on Mobile Systems, Applications and Services (MobiSys)*, 2007, pp. 57–70.
- [27] N. Seyff, F. Graf, and N. Maiden, "Using mobile RE tools to give end-users their own voice," in *Proceedings of the 18th International Requirements Engineering Conference (RE)*. IEEE, 2010, pp. 37–46.
- [28] T. Wehrmaker, S. Gärtner, and K. Schneider, "ConTexter feedback system," in *Proceedings of the 34th International Conference on Software Engineering (ICSE)*. IEEE, 2012, pp. 1459–1460.
- [29] N. Seyff, G. Ollmann, and M. Bortenschlager, "AppEcho: A user-driven, in situ feedback approach for mobile platforms and applications," in *Proceedings of the 1st International Conference on Mobile Software Engineering and Systems (MOBILESOFT)*. New York, NY, USA: ACM, 2014, pp. 99–108.
- [30] D. Wüest, F. Fotrousi, and S. Fricker, "Combining monitoring and autonomous feedback requests to elicit actionable knowledge of system use," in *Proceedings of Requirements Engineering: Foundation for Software Quality (REFSQ)*, E. Knauss and M. Goedicke, Eds. Springer, 2019, pp. 209–225.
- [31] D. Johnson, J. Tizard, D. Damian, K. Blincoe, and T. Clear, "Open CrowdRE challenges in software ecosystems," in *Proceedings of the 4th International Workshop on Crowd-Based Requirements Engineering (CrowdRE)*, 2023, pp. 1–4.
- [32] F. Soares, J. a. Araújo, and F. Wanderley, "VoiceToModel: An approach to generate requirements models from speech recognition mechanisms," in *Proceedings of the 30th Annual ACM Symposium on Applied Computing (SAC)*. New York, NY, USA: ACM, 2015, pp. 1350–1357.
- [33] S. Maghilian and M. R. Kumar, "Sentiment analysis on speaker specific speech data," in *Proceedings of the International Conference on Intelligent Computing and Control (I2C2)*, 2017, pp. 1–5.
- [34] ISO/IEC, "ISO 11783-1:2017(en) (ISOBUS) – Tractors and machinery for agriculture and forestry – Serial control and communications data network – Part 1: General standard for mobile data communication," ISO/IEC, Tech. Rep., 2017. [Online]. Available: <https://www.iso.org/standard/57556.html>