Dahlia AI Tutor Pilot Project

Enhancing the future of student learning

September 16, 2022

Landon Garrison
Founder and CEO
Persona
landon@persona-ai.com
229-116 Research Drive

Executive Summary

Secondary students in Saskatoon and across the province are adjusting back to the normal cycle of in-person classes for the academic year. Unfortunately, many students learning has been negatively affected by the COVID-19 pandemic [2]; Students have been struggling to keep up with the course material and are falling behind [2]. Outside of the pandemic, student enrollment continues to increase which will make it more difficult for students to get one-on-one help from teachers and educational assistants [1][3].

To address this issue, Persona is proposing a pilot program to test the effectiveness of a project we call "Dahlia". Dahlia is a conversational AI system that is able to help students with their courses in Math, Science, History and Computer Science. Dahlia is capable of answering any question a student may have, providing examples and helping students better understand the material. Through the conversational agent, students will be able to engage with the material in a more natural and personalized way.

The first phase of our pilot project will begin in the first semester of the high school year (pending approval) and will continue until the end of the semester. During this time, students will be able to use Dahlia through Facebook Messenger on their smartphones and will be able to use the app as often as they choose. During the duration of the semester, we will use the feedback we get from students and teachers to make improvements to expand our beta testing into the 2023 winter semester. The demographics of students will be dependent on actual enrollment in the classes, but we will aim for an even distribution of strong and weaker students.

We are requesting permission to enroll a small group of students (maximum of 10 students) to use Dahlia per class we partner with. The pilot project will consist of interviews with students every three weeks to measure how they are using the app, what benefits they feel they are getting and where it could be improved. More importantly, we want to see if students will be more willing to engage with educational technology that is highly personable compared to current alternatives. All interviews will be conducted outside of class time on school premises to avoid any interference with classes/instruction.

If the pilot is successful, we believe it will have a positive impact not only on student learning outcomes but also on student engagement. In addition, we believe this software will be helpful to teachers by providing an analytical view of their students and how they are learning. We are excited about the potential of this software and we hope that it has a positive impact on students and teachers.

All costs for the pilot project will be covered by Persona. The schools and the division will not be financially responsible in any way. The only assistance we are requesting is permission to conduct our pilot project in person with the students.

About Persona

Persona is a Saskatoon-based startup that builds Artificial Intelligence systems to understand human language. Our mission is to provide AI systems that positively impact areas of accessibility, education, and the workplace.

Methodology and Pilot Project in Practice

This section provides more detail about the pilot project, hypotheses, methodology for how we will measure the effectiveness. We will also detail any planned data collection and how it will be used.

Thesis, current software and hypothesis

Our main thesis is that current students enrolled in secondary schools are struggling to keep up with the course material and are falling behind as a result of the COVID-19 Pandemic [2][3]. With trends showing that student enrollment will continue to increase over time and budgets decreasing [1], students will have less access to one-on-one help from teachers and educational assistants. We suspect that this will not only have a negative impact on student learning outcomes going forward but that it will also have a negative impact on student engagement in school.

To date, most educational software has been focused on making the teacher's job easier with tools for attendance, grading and scheduling. Unfortunately, there has been very little focus on student learning itself. For applications that do attempt to focus on student learning, they tend to focus on regurgitating knowledge, memorization and/or course delivery. These applications are often based on a multiple-choice format, where questions are answered and graded automatically. The problem with these approaches is that they often don't help students learn and they are often not personalized to the student's needs [4].

To address these issues, we propose an Artificial Intelligence system that students can have a one-to-one conversation about the material they are learning, similar to a teacher; We call this system “Dahlia”. Dahlia will be compatible with any course material being learned and will be able to help students with specific questions they have during their studies. In our first pilots, the focus will be on Math, History, Science and Computer Science.

Our hypothesis is that providing students with an entity that they can study with and ask the questions they have may increase students' understanding of the material being learned. In

addition, we hope to increase student engagement by allowing students to have more meaningful interactions with the material they are learning. Because this interaction is taking place over a conversational interface (e.g. Facebook Messenger), students will be able to have this experience on their own, without needing a teacher or educational assistant present. While our intention is not to replace educators, we do believe there is a major gap once students leave the classroom and do not have access to their instructors [4]. Our goal is to provide a system that can help fill that gap.

Pilot Project and measuring effectiveness

To test the effectiveness of our Dahlia, we propose to conduct a pilot project with a small group of students attending the class in the fall semester. The pilot project will begin in the first semester of 2022 and will continue until the end of the semester. We cannot provide exact dates at this time, as it is dependent on any additional approvals, audits or requests schools have for the pilot project that we may have to work to comply with. During this time, students will be able to use Dahlia through a mobile app on their smartphones. At the end of the semester, we will use the feedback we get from students and teachers to make improvements and then expand our project into the 2023 winter semester.

The main way we will measure the effectiveness of Dahlia is by conducting tri-weekly interviews with selected students. The purpose of the interviews is to get a sense of whether students are using the app, how they are finding their experience, and whether they believe the app is helping them learn better. Interviews with students will determine questions we have going into this project:

1. Can AI tutors help student learning?

2. Will students feel more comfortable asking AI tutors questions compared to a teacher?

3. Will students be more likely to engage with the material if they have an AI tutor available that can help in a personalized manner?

At this stage, the answers to these questions are not obvious and we need to test them through a pilot project. By working closely with students in schools and throughout the semester, we

hope to get better answers to these questions. We are particularly interested in questions two and three; As part of our hypothesis, we speculate that most students (regardless of academic performance) tend to have increased anxiety to ask teachers for one-on-one help (this however will also need to be proven through our pilot). A portion of this pilot project will be tracking to see if students feel more comfortable talking with Dahlia and whether they are more likely to engage with the material as a result of using our software.

A separate question we also have is how students who do not speak fluent English would use a tool like Dahlia. For the purposes of phase 1 of our pilot (detailed later), we will only be testing our software with students who are fluent in English; Later in the year however (phase 2), we plan to expand usage to students whose first language is not English. As part of our hypotheses, we speculate that non-English speaking students would benefit from our app as we are developing capabilities for Dahlia to help students in their desired language (i.e. Dahlia can talk to a student in Mandarin). To reiterate, this hypothesis needs to be tested further with student participation.

Target courses

For phases 1 and 2 of our pilot project, Dahlia will be capable of helping students in Math, Science and History. Specifically, we intend to target students enrolled in the following subjects:

- Math 9, Pre-Calculus 10, Pre-Calculus 20, Calculus 30
- Science 9, Science 10, Biology 30
- History 20
- Computer Science 30 or lower related courses

If we decide to add any courses to our targeted list, we will notify the schools and teachers we are working with.

An important thing to note is we expect our software won't be perfect in every subject of the pilot project. In early testing, we have discovered areas where it makes mistakes such as repeating itself, incorrectly computing numbers and not explaining concepts in high enough detail. These issues however are predictable to see when they might occur, so we can instruct students and teachers to be mindful of these scenarios.

Privacy and data collection

Because this project will be working with students in a school division, there may be some privacy concerns around data collection. The only personally identifiable information (PII) Dahlia uses is the student's email associated with their Facebook account log in to the app. Outside of collecting data from interviews, we plan to collect the following metrics from our app:

1. How often Dahlia is used by the students
    - This is a signal to see if students are finding the app useful.

2. Sentiment analysis of conversations
    - We use sentiment analysis to determine if conversations with Dahlia are positive, neutral or negative in tone. We automatically measure this by using separate AI models to provide classifications on conversations.

3. Tracking conversation types
    - While we assume most conversations with Dahlia will revolve around school material, Dahlia is capable of supporting students when they are discouraged and need support. We want to track to see if students engage with Dahlia to seek support and if so, we will expand our software to provide additional content in the future.
    - We will also be collecting the actual interactions students have with Dahlia in the form of a chat log (history). These logs are fully anonymous and cannot be traced back to the student. We use this information to train newer versions of our models (eg. make improvements).

Transparency for how we anonymize data

To ensure trust between the school divisions, students and parents, this sub-section will provide exact details for how we anonymize data collected in the app.

Because the core of our app is based on conversations with Dahlia, there is more room for personal information to be exposed compared to a traditional app (eg. Desmos). To provide a private experience where we can also collect the necessary analytics we need, we “encode” student accounts. This means that anytime a student interacts with our app, their account

interactions will only be represented with a series of randomly generated numbers. **We cannot decode, decrypt or decipher these numbers because these codes are auto-generated from Facebook**. In other words, **the only possible way to understand who is behind the conversation with Dahlia would be to have access to Facebook's source code, which we do not have**. The conversations however are left in their original format; This is a critical piece to understanding how our app is being used by the students.

Transparency and concerns around using social media as our primary platform

We have decided to integrate Dahlia into Facebook messenger as our prototype. Parents, schools and school divisions may have concerns or questions about why we decided to do this, and if it exposes students to any additional data collection.

Why we decided to use Facebook

Our decision to integrate Dahlia into Facebook Messenger is for student ease of access. Currently, Facebook has approximately 3 billion users and 1.9 billion daily active users. Because of this, the probability that students we work with will have access to Facebook is quite high, which is why we decided to use Facebook as our primary platform. Facebook messenger in particular has a fluid experience that we think most students will be familiar with, which was another ease-of-use decision.

Social media data collection

Our decision to integrate with a social media platform may raise some concerns around student/minor data collection, which we fully understand. This section will provide complete transparency on this portion of the app.

Facebook provides Persona and many other companies around the world with a "Developer API" to its platform. An API (stands for Application Programming Interface) is a tool that allows two pieces of software to communicate with each other, by making a request to a server and the server returns information to an application. For example, Dahlia (our software) communicates with Facebook Messenger (their software) to send or read messages that students send to Dahlia.

The reason this is mentioned is that Facebook requires every app they use to go through a stringent “app review” process before requesting new functionality, including user data that could be more sensitive. This means that every app that uses Facebook’s developer services must have a legitimate and authentic reason for using its APIs. So for example, Persona has to request permission from Facebook to send students messages via Messenger (eg. Dahlia sends students messages), which was approved because it is a core functionality of our app (i.e, we need to be able to reply to students who message Dahlia). If for example, we were to request more sensitive information (like a user's birthdate or age), we would need to prove that we have a legitimate reason for requesting information from Facebook.

The specific APIs we use from Facebook that we have been approved for are the following:

API Title	Facebook Description	Our Reason for Requesting	Example usage in practice
pages_messaging	The pages_messaging permission allows your app to manage and access Page conversations in Messenger. The allowed usage for this permission is to create user-initiated interactive experiences, send customer support messages or confirm bookings or purchases and orders.	We use this API to read and respond to student messages that they send to Dahlia.	Our core functionality with our app is when students talk with Dahlia
email	The email permission allows your app to read a person's primary email address. The allowed usage for this permission is to let end users log into your app with the email address associated with their Facebook profile. This permission is automatically granted to all apps.	We use the students/users email to authenticate them for the first time. This prevents our app from receiving spam and unwanted usage.	When the students use our app for the first time, they must select the email associated with their Facebook account.

public_profile	The public_profile permission allows apps to read the Default Public Profile Fields on a User node. The allowed usage for this permission is to authenticate app users and provide them with a personalized in-app experience. This permission is automatically granted to all apps.	We do not use this API, but it is granted to every app built on top of the Facebook developer platform.	Not used in practice.

Safeguards and protections

Conversational AI software pose unique challenges compared to other types of educational software. Because our system is based on conversation, there is a greater risk that students will ask inappropriate questions to Dahlia or prompt it to generate incorrect responses. To address these areas, we have employed a "content filter". A content filter is a separate AI program that acts as a "second set of eyes" to ensure that generated responses are appropriate, non-harmful and non-toxic. Employing content filters is a common practice in this type of software across the industry.

Like any AI software, occasionally our content filter will make a misprediction. However, we are confident that our content filter will catch harmful responses with high accuracy and if failures occur, more often it will be "false-positive" predictions. In the context of a content filter, this means that it may redact a response that was not actually harmful or toxic. If a student believes that an inappropriate response was given by Dahlia, they will be able to flag it in our app or by communicating it directly with us in our interviews.

Our content filter currently categorizes text (replies from Dahlia) in the following ways:

1. **Safe:** The reply is safe to return to the student and contains nothing inappropriate, harmful or toxic

2. **Sensitive:** The reply contains sensitive material (ex. politics, religion, etc.) that may be offensive to some users.

3. **Unsafe:** The reply contains inappropriate, harmful or toxic material. The response will be redacted and re-generated.

A weakness we have identified in development is our content filter can be overly cautious about historical topics. Because topics in history (ex. World war 2, the holocaust, historical genocides) can be sensitive in nature, our content filter may redact these topics, even if they are being discussed in a non-harmful way. To address this, we have started the development of a "context-aware" content filter, which takes into consideration the conversation and topics being discussed. This is however a very complex and nuanced problem and is currently being worked on by most major companies working in this space.

The next page shows a visual diagram of how our content filter works.

"Virtual Tutor" = "Dahlia", updated 12/11/22

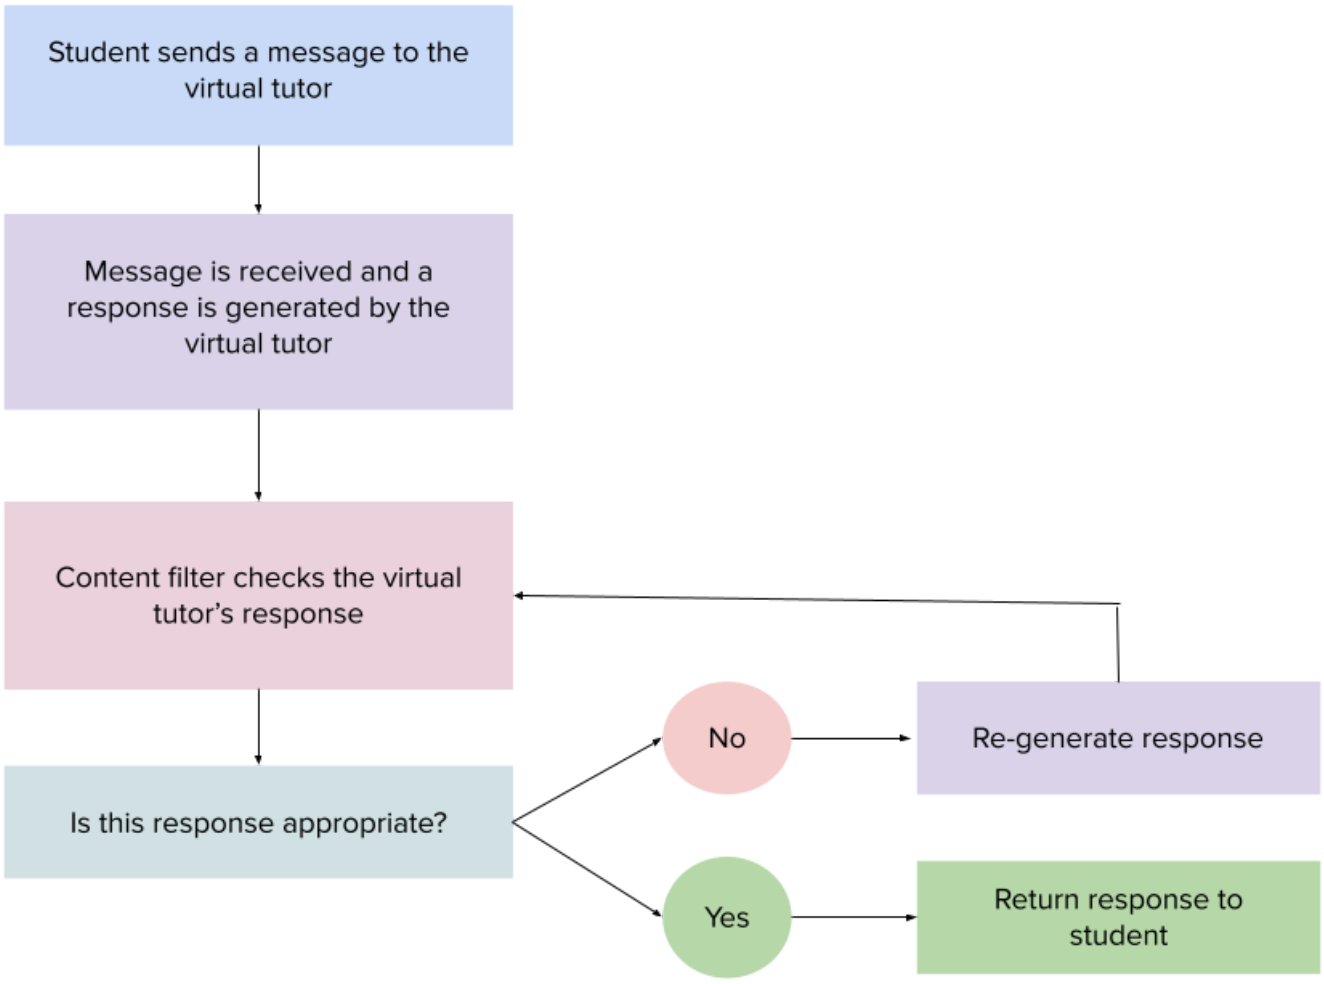

Assumptions and Expectations

This section will summarize the assumptions and expectations we have for the pilot project.

1. Demographics of students

We have made no assumptions about students who will be enrolled in our pilot project. Our only expectation is that students using our app must be committed to providing feedback and using the app.

2. Access to technology

Our software is available in the form of a mobile app; Students will need access to a smartphone with internet access. We have assumed that the majority of students enrolled in schools will have access to these technologies.

3. Students with social media accounts

We have integrated our app with Facebook Messenger. The reason for this is that we assume most students enrolled in secondary schools will have social media profiles. Integrating our app into these platforms will make for a highly frictionless experience as no additional accounts will be required.

4. In-person classes and interviews

We have assumed that most students will be in person at their school during the fall semester. Pending approval, we have assumed that we will be allowed to conduct interviews with students in person at the school. If either of these assumptions does not pan out, we will need to adapt and make accommodations for students. To re-iterate, all interviews will be conducted outside of class time at the student's classroom.

5. Students using the software speak fluent English (Phase 1)

We have assumed that all students enrolled who participate in the first phase of the pilot project are fluent in English. Currently, our software currently only supports conversations in the English language. We are working on expanding our software to support more languages in the future, but at this time, we will only be able to support English-speaking students. If however, we have a majority of students are not native English speakers, we have measures to add automatic translations into our app. We plan on having multi-lingual features ready by phase 2 of our project.

6. All costs for teachers, educational assistants, school employees and school programs will be covered by the School Division

Persona will not be responsible for incurring any of the costs associated with salaries for Division staff, employees or any faculty. We will only be responsible for covering the costs of our pilot project, software development, hiring of research assistants and the hosting of our app.

7. Autonomy to run the pilot project

We are under the expectation that the School Division and/or school will have no involvement with the pilot project and the beta testing itself. We understand that there may be requests from the schools and/or the division we work with to adjust portions of our project, which we are happy to comply with and take suggestions.

8. Data collected will be kept private

We are under the expectation that any data collected from students and teachers will be kept private and will not be disclosed to anyone outside of Persona, including the Schools and the Divisions for the duration of the pilot project. No data will be published or sold as our only intention is to gather a better understanding of how Dahlia is used in a real-world setting.

9. Any data collected will remain the property of Persona

We are under the assumption that any data collected as part of this pilot project will remain the property of Persona. We are open to sharing our findings with school divisions and plan to collaborate heavily with divisions.

10. Research participants are volunteers

All students and teachers participating in the pilot project will be volunteers. We will not be paying any students or teachers to participate in the pilot project.

11. Concerns from the Schools or Division will be addressed by Persona

We are under the expectation that Persona will be responsible for addressing any concerns the schools or divisions may have about this project or the app.

12. Informed consent and student privacy

We are under the expectation that we will be able to conduct interviews with students if they are willing to participate. We are also under the expectation that we will be able to collect data on student usage of our app. In both cases, we will ask for written consent from the student and parent (if the student is under 18). Our app will never ask for any personal information and all conversations with Dahlia are fully anonymous.

13. No payments for the pilot project will be made by the school division or the school

We are under the expectation that no payments will be made to Persona for the pilot project. If discussions of payments arise (i.e to license our software), they will be separate from this pilot project proposal.

14. Dahlia and additional software are property of Persona

We are under the expectation that all AI models and technology developed as part of this pilot project are the property of Persona and that we will not be disclosing the intellectual property to the school divisions or the schools. We also reserve the right to license this IP to other companies or school divisions in the future. If there is the disclosure of any code, proprietary information or trade secrets, it will be released under a Non-disclosure agreement.

15. Outside of authentication, no will be collected from students' social media profiles

As mentioned, we have integrated Dahlia into Facebook messenger for convenience to the student. No data will be collected from students' social media profiles or their messages outside of the ones sent to Dahlia (i.e, when students use the app). We cannot “scrape” (collect) any other data from students' profiles as Meta (Facebook) does not allow this access without explicitly requested permission.

17. Significant changes to our plan

If we make significant changes to our research plans, we will notify the schools/school division by email, phone or another preferred method of communication.

18. Students will not sell or give other students access to their account

Dahlia is a “beta” software, meaning it is still in research and development. We prohibit students from allowing other students who have not been selected in our pilot to use our app.

19. People conducting in-person interviews will have criminal record checks

To ensure safety, all people conducting research will be required to complete a criminal record check prior to meeting any students or staff in person or online. We would appreciate guidance from the schools we work with if there is a specific type of criminal record check that would be preferred. **If no guidance is provided from schools or school divisions, we will assume the correct criminal record check is the “Record Check without the Vulnerable Sector” variant.**

20. Our software is not perfect

We do not expect Dahlia to be perfect at all times and we expect it to make mistakes. While we believe we have reached a point where these mistakes are not common, we will not attempt to over-promise students and staff that it is without flaws. As stated earlier, the portion of this pilot project is to understand where it fails. We will communicate with staff and students about what situations will fail in and be honest about current capabilities.

Outcomes of the Pilot Project

This section will go into detail on what we hope to accomplish with the pilot project, its importance and how we will measure success.

Goals and measuring success

The main goal we hope to achieve through this pilot project is to determine whether we can positively affect a majority of the students learning outcomes that use our software. We define a successful learning outcome as a student who has found our app helpful, engaging and inclined to use our app for other courses. Engagement is a metric we are most concerned with (measured in time used using the app) because we are still unsure of how students will use our app.

A second major goal of the pilot project is to see if our software can provide positive learning outcomes for students who aren't academically inclined or who have struggled in the past. This is a major challenge for any educational software, but we feel strongly that our approach provides unique capabilities that have largely been unexplored. A major success for the pilot project would be if our software could provide motivation for struggling or unenthusiastic students to engage and learn their course material.

While our primary focus is on student learning, we also hope to provide value to the teachers. As part of our goals, we hope that teachers find our software helpful and that it makes their job not only easier in some areas, but provides better insights into their students. While this pilot project will not engage in direct analytics for teachers (extracted from the app), eventually we plan on using our learnings from this pilot project to also help teachers with AI software.

Importance of the pilot project

The pilot project is being carried out to test the effectiveness of our software and determine whether it is ready to be scaled up to a larger population of students. We believe that our software could help a large number of students, but we need to test it with a small group of

students first. The pilot project will be a crucial step in allowing us to see if our vision is correct and whether students are finding it helpful and useful in their learning.

We believe that this software has the power to positively impact a large number of students and ultimately, the future of educational technology. By testing this software with a small group of students first, we will be able to make large-scale improvements in future versions of our software. At Persona, we view this pilot project as a pivotal point in educational software.

Deliverables

This section will go into detail on what we expect to deliver as a result of the pilot project.

Deliverables for students

The main deliverable for students is that they have a positive experience using our software and have a better understanding of the material they are learning. This is the most important area that we expect to deliver. As a result of the pilot project, we will also be making adjustments to the software and improving it based on feedback from the students and teachers. While we hope to deliver a positive experience to all students, we are prepared for some students to have less-than-perfect experiences. While this is expected, we will work with these students to understand how we can improve our software for them.

As mentioned earlier in the hypothesis section, we hope to help students have access to affordable assistive technology for their courses in school. Based on conversations had with teachers, we have a strong conviction that students would benefit from technology that they can ask for help at any time of the day in a judgment-free environment.

Deliverables for Persona

As a result of the pilot project, we will be delivering a tangible and analytical view of whether our software is being used and if it is helpful to students. We will be making adjustments to our software based on feedback and will be preparing to scale it up to a larger number of students. We will also be delivering a report on the pilot project and the effectiveness of our software.

Deliverables for teachers and school divisions

As mentioned earlier, we view this pilot project and a pivotal point in providing the Saskatchewan and Canadian educational systems with truly helpful technology. We think there is a potential to deliver a highly innovative product in the future for schools and school divisions.

Findings in this first pilot project could shape the next major milestone in the Saskatchewan education system, and we would be proud to deliver a product to the schools we work with in this initial pilot project.

Timeline

This section will go into more detail about the timeline of the project and what we hope to accomplish at each stage.

Phase 1 (~ September or October, pending approvals)
We hope to begin the first phase of the pilot project will begin in the first semester of the 2022 year and will continue until the end of the semester. During this time, students will be able to use the Dahlia through their smartphones. At the end of the semester, we will use the feedback we get from students and teachers to make improvements and then expand our testing into the 2023 winter semester. The purpose of phase 1 is to get a sense of the utility and effectiveness of our software. We will be able to determine if the software is effective by measuring student learning outcomes and by conducting interviews with students.

Phase 1 of our pilot project will consist of tri-weekly interviews with students after their class. At the end of the semester, we will use the feedback and data we have collected to improve the software and make any required changes. This will also serve as a re-evaluation point to see if we will move into the second phase of pilot projects in the winter semester.

Phase 2 (January to June)
The second phase of the pilot project will begin in January 2023 and will continue until the end of the school year. The key outcome of phase 2 is to gain a greater statistical significance to measure learning outcomes and how our app is used. Similar to phase 1, we will use the feedback we get from students and teachers to make improvements. At the end of the school year, we will analyze the data that we have collected and determine whether our software is ready to be rolled out to a wider population of students.

Phase 2 will conclude our pilot project and we will be removing student access to the software. At this point, we will determine if the pilot was a success and if we can move forward with our software.

Sign-off and acceptance

Please note: This sign-off is not intended for parents. Please see our "parental consent form" for parent consent.

We are excited to start our pilot project. We are asking for written consent (via email or mail) and a signed copy of this form.

I _____ have read the proposal and understand the purpose, scope and tasks involved in the Dahlia pilot project and give my written consent to Persona to proceed with the pilot project for:

Phase 1: Fall Semester (Date is pending approval)
Phase 2: Winter Semester (January 2023 to June 2023)

x_____

Signature

_____

Representative title (i.e: Superintendent/principal)

___ ___ / ___ ___ / __ __ __ __

D D M M Y Y Y Y

References

[1] *Saskatoon Public School Division Budget and Report, 2022* https://www.spsd.sk.ca/division/reportsandpublications/Documents/SPS%20Budget%20Report%202021-22.pdf

[2] Kuhfeld, Soland and Lewis, *Test Score Patterns Across Three COVID-19-impacted School Years,* https://edworkingpapers.com/sites/default/files/ai22-521.pdf

[3] Pennino et.al, *Student Anxiety and Engagement with Online Instruction across Two Semesters of COVID-19 Disruptions,* https://journals.asm.org/doi/epub/10.1128/jmbe.00261-21

[4] Minkos & Gelbar, *Considerations for educators in supporting student learning in the midst of COVID-19*, https://onlinelibrary.wiley.com/doi/epdf/10.1002/pits.22454