

Intro & Agenda

Aaryan Handa

SP26-54

S.E.A.T. Phase 2 Expansion Plan: Transforming a Smart Chair MVP into an Adaptive Ergonomic Intelligence Platform

Agenda

- Introduction and Project Background
- Strategic Overview of the Phase 2 Expansion Plan
- Phase 1 MVP Recap: What was built, technical architecture, results, gaps, KPI's and metrics, potential enhancements
- Phase 2: Strategic Vision
- Phase 2: Technical, AI, and Data Enhancements
- User Experience, Safety, Privacy, and Human Factors
- Business Model, ICP, and Go-To-Market Strategy
- Manufacturing, Cost, and Scalability Plan
- Roadmap + Future of this
- Final Recommendation and Conclusion
- References and Citations

Introduction and Project Background

Introduction and Project Background

General gist

- **The real-world issue:** The dangers of bad posture
- **Why current solutions are limited:** Stagnant and uninnovative market
- **Why S.E.A.T. matters:** Leverage AI in a smart way for the greatest societal impact

Poor posture and prolonged sitting have become much more dangerous, and major issues have been linked to them. Students spend hundreds of hours in labs and classrooms, and employees spend multiple hours staring at a screen on their desks. Remote workers do not realize the long-term implications of constant sitting until they feel a sharp chronic pain in their lower back, “one random day”. Low back pain is important because it is one of the most common musculoskeletal conditions globally, and the World Health Organization even identifies low back pain as the leading cause of disability worldwide. The WHO also estimated that approximately 619 million people experienced low back pain in 2020, with cases expected to rise to 843 million by 2050.¹

Traditional ergonomic chairs address this issue by providing users with better support through features such as adjustable height, lumbar support, armrests, and even manual seat positioning. Traditional chairs provide support once the user has become reactive to bad posture. This creates a gap between users having these features and accurately using them over time. For example, a person may begin a study or work session sitting properly, but after 20 or 30 minutes, they may gradually lean forward, slouch, or shift weight unevenly without noticing. This is the active gap in modern-day traditional chairs. They provide good posture support, but do not mitigate bad posture if the user’s posture becomes uneven or reverts to being bad.

S.E.A.T was created keeping these issues in mind. Smart Ergonomic Adaptive Technology, which abbreviates to S.E.A.T., treats the chair as more than just furniture. It turns it into an intelligent system that can sense, think, adjust, and give feedback.. The core concept of this was to build a smart chair that can analyze the sitting mechanisms and posture of the user, then classify the user’s posture through machine learning, followed by physically adjusting parts of the chair using motorized mechanisms, and provide posture analytics through a connected application. On a simpler and foundational level, the chair is designed to move from passive support to active posture correction. The motivation and purpose behind a product like this was to shift the chair from being reactive to proactive.


The ergonomic chair market is a rapidly growing field, and the market demand has grown significantly over the last few years. The growth rate is projected to be around 8.5%, rising from 12 to 25 billion by 2034². There is a huge demand for this product, and while that does not directly correlate with the automatic success for a product like this, it gives individuals a great opportunity to look at a fast-growing niche market.

Another major reason why S.E.A.T. matters is that, not only does it look at posture as a physical issue, but also as a behavioral issue. A lot of people know that they should sit straight, but maintaining that posture throughout the day is where the real problem begins. Most people do not actively think about their posture while doing homework, coding, studying, working, or attending meetings. This is where the product becomes more than just an ergonomic chair. It creates a feedback loop where the chair is not only supporting the user, but also learning from the user and helping them become more aware of their own habits over time. That is what makes the idea more powerful than just adding more cushion or another manual lever to a regular chair.

From an engineering and building standpoint, the purpose of S.E.A.T. is also important because it combines multiple fields into one product. It includes hardware systems like pressure sensors and motors, software usage through the user application, and artificial intelligence through the machine learning model that classifies posture and does the actual thinking. This type of product is valuable because it shows how AI can be used in a practical and physical way, instead of only being used for chatbots, automation, or data analysis. In this case, AI is being used to solve a real human problem by improving comfort, health awareness, and long-term sitting behavior, which has a direct impact on human health. That gives the project a stronger societal purpose because almost everyone sits for long periods, but very few people have a system actively helping them correct their posture in real time.

Overall, the demand behind S.E.A.T. is strong because the problem is a major issue and has serious health implications, the market is rapidly growing, and current solutions still leave a major gap. Phase 1 did the foundational work by proving that this idea can actually function as an MVP through pressure sensing, machine learning classification, motorized adjustment, and user analytics. However, Phase 2 is where the product can move beyond being just a working prototype and become something that has real demand in a rapidly growing multi-billion-dollar industry. The main idea of this report is to show how S.E.A.T. can become 10-fold more transformative, from a basic smart chair into something safer, smarter, and more scalable, all while being proactive instead of reactive, and a more commercially valuable ergonomic intelligence platform. This includes improving the technical system, making the AI more predictive and personalized, strengthening safety and user experience, creating a clearer business model, and building a realistic path for manufacturing, validation, and market growth, all while holding an incredibly holistic approach.

Global Ergonomic Chair Market Size, 2021-2034 (USD Billion)



Graph depicts the growth for this industry

Strategic Overview

Strategic Overview and Phase 2 Expansion

General gist

- **Phase 2:** Turning the working MVP into a serious product direction
- **Transformation:** Turning a reactive MVP into a proactive product
- **Impact:** Phase 2 is where the chair becomes safer, smarter, more personalized, more scalable, and more valuable in a real market

Phase 1 of S.E.A.T showed the general audience what a working MVP looks like. The team was able to create a chair that collects pressure data through software systems, uses machine learning to classify bad posture, physically adjusts through hardware systems, and displays basic posture information through an application. That is a strong foundation because it shows that the project was turned into a working concept instead of just ideation. It is a real MVP that connects hardware, software, AI, and mechanical movement into one system. However, Phase 1 should still be viewed as the starting point, not the final product. The main purpose of Phase 2 is to take that working foundation and turn it into something more intelligent, more reliable, and more realistic for actual users.

The biggest shift in Phase 2 is moving S.E.A.T. from being **reactive** to being **proactive**. Phase 1 shows a reactive approach where the chair is illustrated to move hardware systems to fix the user's posture once it is bad. That is still useful, but it does not fully solve the problem because posture usually gets worse gradually over time for most people. It is extremely rare for a person to go from a perfectly normal posture to one that is slouching or leaning on either side very drastically and suddenly. The change is often very slow and, in most cases, unnoticed. This is why Phase 2 should focus on predicting posture decline earlier instead of only reacting once the posture is already bad. This matters because ergonomics is not just about comfort. NIOSH explains that the goal of ergonomics is to prevent discomfort and injuries at work, and OSHA also describes ergonomics as fitting the job to the person in order to reduce muscle fatigue and work-related musculoskeletal disorders³.

This is where S.E.A.T. has a huge market opportunity because the foundational idea is to react instantly rather than gradually. A regular ergonomic chair gives support, but it does not understand the user. It does not understand the behavioral habits and postures of a user and fails to see the user as a person rather than a use case. The Phase 1 prototype can sense and respond, which has moved the chair in the right direction. Phase 2 should push it further by making it an **ergonomic intelligence platform**. That means the chair should not only know what posture the user is in at one moment but also understand patterns over time, learn the user's normal sitting behavior, recognize when posture is starting to decline, choose the least invasive correction, and help the user build better sitting habits. On a simpler level, Phase 2 should make the chair feel less like a machine that suddenly moves and more like a smart support system that learns how to help each person on a more subjective note.

A huge part of this expansion and enhancement is improving the AI and data strategy. The current version of the chair uses machine learning to understand posture based on sensor readings, but Phase 2 should make the model more predictive and personalized. To put it in simpler terms, the goal of phase 2 should go from “the user’s posture has become bad and I need to fix it” to “the user’s posture is gradually regressing and based on predictive patterns, I need to take smart proactive action”. This is a more advanced way of thinking because it treats posture as a behavior over time, not just a single label. Random Forest was a very solid starting point for this type of project because it combines multiple decision trees and averages results to improve predictive accuracy and reduce overfitting, but Phase 2 should go further by improving validation, tracking model confidence, and collecting more diverse data across all different types of test cases⁴.

Phase 2 also needs to make the chair feel safer and more comfortable. Since S.E.A.T. is designed to physically move a person sitting on it, there is no room for error. The chair should not move every time the model makes a prediction, because most people constantly move a bit while sitting. Not every movement is bad posture, and not every bad posture is held for long periods of time. A person leaning for two seconds is normal. A person leaning for twenty minutes may be a problem. This is why Phase 2 will include confidence-based and smart correction logic and false-correction prevention. The system should only physically adjust when it is confident that the user is in a sustained poor posture pattern instead of abrupt movement. If the model is only somewhat confident, the chair can start with a softer intervention, such as a visual reminder, app notification, or small coaching message, before moving into physical adjustment.

Another important enhancement of Phase 2 is personalization. Every user sits differently. A taller person, a shorter person, a heavier person, a lighter person, an athlete, a student, or even an office worker may all create different pressure patterns on the same chair. Even if two people are both technically sitting upright, their sensor readings may not look the same due to different factors like weight distribution. Phase 2 will build on the original calibration idea and turn it into a deeper user profile for the sake of personalization. The chair should learn each user’s default posture, common posture habits, correction preferences, and even a small outlook on fatigue patterns. This would make S.E.A.T. more accurate over time because it would not treat every person as the same user. It would begin to understand what “good posture” and “bad posture” look like for each different individual.

The expansion plan also needs to be incredibly transformative. With the advancements today, what separates an MVP from a multi-billion-dollar product is the idea of edge computing⁵. A working prototype can depend on a host computer, exposed wiring, and prototype-level parts, but a real product cannot stay that way. The sustainability, reliability, and, to put it bluntly, the “lack of advancement” is where Phase 2 should move toward edge-based processing, where the chair can run posture detection locally through an onboard computing device instead of needing a laptop nearby. The usage of cleaner hardware design, custom PCBs, stronger mechanical components, enclosed wiring, and modular parts that can be repaired or replaced more easily are essential to turn this from a basic MVP to a scalable product. This is important

because a chair that works once in a demo is not the same as a chair that works every day for millions of users.

To summarize, the Phase 2 expansion plan is about turning S.E.A.T. from a working MVP into a serious product direction. Phase 1 answered multiple questions: **Can the chair sense posture, classify it, adjust physically, and show user feedback?** Phase 2 now has to answer the bigger question: **Can this become reliable, safe, personalized, scalable, and valuable enough for real-world demand?** The rest of this report will focus on how to make that happen through stronger AI, better hardware, safer motor control, improved user experience, privacy-first analytics, a clearer business model, and a realistic plan for validation and market growth. This is where S.E.A.T. can move from being a basic smart chair prototype into a product that has real potential in a growing multi-billion-dollar industry.

Phase 1: Reactive MVP	Phase 2: Proactive Platform
The user has shown bad posture	The user's posture is slowly declining
Sensors detect pressure	Sensors track pressure over 'x' time
ML classifies posture decline	AI is used to predict over time the posture decline
Motors adjust the chair	The system chooses the least invasive intervention method
App displays feedback	The app is used to show the user their body patterns to actively work towards better posture

The table depicts the ideology behind Phase 1 vs Phase 2

Phase 1 Recap

Phase 1 recap

General gist

- **Phase 1 Technicalities:** The technical working side of the functioning MVP
- **Business Model and Positioning:** The current business model plan
- **Metrics, Enhancements, and the forward path:** The transition from Phase 1 to 2

The main goal of Phase 1 and this capstone project was to prove that a working MVP could be made with the maximized use of AI, hardware, and software systems. Instead of only creating a concept or a basic ergonomic chair design, the team built an MVP that connected hardware, software, machine learning, and mechanical movement into one complete product loop. On a simpler level, the chair was designed to understand how the user was sitting, classify that posture, physically adjust the chair to fix the posture, and then show feedback through an application. This was critical as a foundation because it proved that multiple systems could be used to work together to make an innovative product, which can have real-world use cases.

To better understand this product, we need to dive deeper into the hardware, software, and AI usage. The chair used 10 pressure sensors across different parts of the seat and backrest to understand how the user's body weight was distributed. A simpler way to think of this is that these sensors acted almost like small weighing points across the chair. If a user sat upright, the pressure pattern would look more balanced. If they leaned left, leaned right, slouched, or sat without back support, the pattern and the weight distribution would shift. That sensor data was then read through the Arduino-based system and sent forward as numerical values that the machine learning model could understand and analyze⁶. This is where the chair started moving from being a normal furniture product into a smart sensing system.

Then came the usage of AI, which was used for classification. The system used pressure sensor readings as inputs and then classified the user's posture into specific categories. Two different machine learning models were used. The SVM(support vector machine) and the Random Forest models were used, and the random forest model gained a higher accuracy and was used as the AI model for the creation of this MVP. This made sense because Random Forest combines multiple decision trees and averages their results, which can improve predictive accuracy and massively reduce overfitting when used properly.

Once the AI model was able to predict the posture, the system moved into the correction layer. This was the selling point and the main idea behind this MVP, as the correction model is essentially the root of this idea. The system used motors, motor drivers, gears, and mechanical sliding components to adjust parts of the chair. In simpler terms, the model would identify a posture issue, the system would decide what movement was needed, and then the motors would adjust the chair to support correction. This is what made the MVP into something of substance because the motors would act, making the chair into a real functioning MVP rather than an ideated product.

The software aspect showed the user what the chair was actually doing. The application displayed posture data, live sensor feedback, heatmap-style visuals, trend tracking, session history, and calibration features. This mattered on a very crucial level because bad posture is more of a mental issue rather than a physical one⁷. Most people know they should sit properly, but they do not actively think about posture while doing their everyday errands and living a routine life. The app helped turn invisible sitting habits into visible feedback. That gave the product more value because the chair was not only correcting posture in the moment; it was also helping the user understand their patterns over time.

From a business model standpoint, Phase 1 had a clear starting idea in terms of technicalities but not yet a fully developed strategy. The end goal of every MVP is to create a product that can ultimately sell, which requires the business model to be as essential as the engineering design. The product was positioned as a more affordable, adaptive, and intelligent alternative to traditional ergonomic chairs. That direction made sense because the ergonomic chair market is already large and growing. Fortune Business Insights estimates the global ergonomic chair market at \$12.40 billion in 2025 and projects it to reach \$25.12 billion by 2034⁸. This supports the idea that there is real demand for ergonomic seating solutions. However, Phase 1 still left massive room to sharpen the business side of things. The product was described as useful for schools, workplaces, homes, and even healthcare settings (rehabilitation centres), but a real go-to-market strategy cannot target everyone at once.

The business direction for Phase 2 will be 10-fold more extensive and focused. If a product cannot hold any sort of monetary value and ultimately sell and grow, it is useless in a real-world setting. Instead of selling S.E.A.T. broadly as a smart chair for anyone who sits, the product should start with high-density environments where people sit for long periods and where one buyer can purchase multiple units. The **ICP** includes universities, computer labs, libraries, corporate offices, call centers, and rehabilitation clinics later on. The early selling strategy should not just be “buy our chair.” It should be framed as a posture wellness pilot, where organizations test the chair for 60 to 90 days and measure posture improvement, user comfort, adoption, and reliability. That creates a much stronger path from prototype to market. The consumer market needs to have high confidence in this product, so that their employees and “use cases” will benefit from the incorporation of this product.


When looking at metrics, Phase 1 proved that the most crucial north star metric is technical reliance. Not only did it point out technical feasibility, to be incredibly important, but it also showed that the MVP excelled in that metric. The MVP showed that the chair could analyze posture, classify it, respond physically, and display feedback. However, Phase 2 needs to move beyond basic feasibility metrics and define more serious product KPIs. These should include posture classification accuracy, false correction rate, model confidence, sensor consistency, motor response time, user comfort rating, calibration completion rate, daily usage, posture improvement over time, and pilot conversion rate. These metrics matter especially because a product that works once during a demo still needs to prove that it can work reliably for many millions of users over long periods of time.

The gaps from Phase 1 aligned perfectly for Phase 2 to take this up a notch. The current MVP is still mostly reactive, which means it responds after poor posture is already detected. The next version(Phase 2) should become more proactive by predicting posture drift before the user fully slouches or leans into a poor postural position. Additionally, the model also needs more diverse data and stronger validation because a high accuracy score on a limited/biased dataset does not guarantee strong real-world performance. The hardware has room for refinement because prototype-level wiring, external computing, and early mechanical parts are not enough for a scalable product. The correction system also needs to become smoother and safer, especially because the chair physically moves while a user is sitting on it, which can create a huge risk of physical lawsuits.

This is where Phase 2 can really take the project up a few levels. The current MVP already proves that the chair can work, but the next step is making it smarter, safer, and more realistic for everyday use. Some of the biggest improvements would be things like predicting posture decline before it gets too bad, using confidence-based correction so the chair does not move for every small shift, adding closed-loop motor feedback, making the adjustments smoother and more comfortable, moving toward edge-based processing, and building stronger personalization and safety controls. On the business side, Phase 2 also needs a clearer idea of who this product is actually being sold to, how it would make money beyond just selling the chair once, and how it could be tested in places like universities, offices, or computer labs. An anonymized dashboard for institutions could also be useful, as long as it focuses on overall ergonomic trends and not individual surveillance. This direction also connects well with NIOSH's view of ergonomics, which focuses on preventing discomfort and injury by designing work environments around people. In other words, Phase 2 should not just make S.E.A.T. more advanced, but make it more useful and realistic for real people using it every day.

Overall, Phase 1 is an incredible start to a successful foundation. It proved that the chair could collect data, classify posture, adjust physically, and show feedback through an application. That is exactly what a strong MVP is supposed to do. However, the real opportunity is not stopping at the MVP. Phase 2 needs to make the product more predictive, more personal, more comfortable, more scalable, and more commercially serious. In other words, Phase 1 proved that S.E.A.T. can work. Phase 2 needs to prove that S.E.A.T. can work well for real users, in real environments, with enough value to compete in a growing ergonomic market.

S.E.A.T. PROJECT PHASE 1: MVP WORKFLOW PROOF OF CONCEPT



Workflow for Phase 1 MVP

Image created using nano banana → Gemini

Phase 2 - Strategic Vision

Phase 2: The strategic vision

General gist

- **What does phase 2 mean?:** Turning the MVP into a product
- **Vision:** The transformation from furniture to a smart AI-leveraged product
- **Impact:** Create a product with meaningful societal impact

Phase 2 could be viewed as the turning point for S.E.A.T. Phase 1 proved that the chair could work as a functioning MVP, but Phase 2 is where the project needs to become more refined, more reliable, and sustainable, and ultimately more realistic to actually sell it. The first version showed that pressure sensors, machine learning, motorized adjustment, and a user application could all work together and create a strong MVP. That is a strong foundation, but still only the foundation. The next stage should not just be about adding a few extra features. It should be about transforming S.E.A.T. from a reactive smart chair into a proactive ergonomic intelligence platform.

The biggest strategic shift for phase 2 is moving from “fixing posture after it becomes bad” to “preventing posture decline before it becomes a bigger issue.” This matters because poor posture usually does not happen instantly or drastically. A person might start a study session sitting upright, but after 20 or 30 minutes, they may slowly lean forward, lose back support, or shift weight unevenly without realizing it⁹. A reactive system can help once the problem is already happening, but a proactive system can catch the pattern earlier, especially if the system has been personalized. OSHA went into great detail with a study that explains that neutral body positioning reduces stress and strain on the muscles, tendons, and skeletal system, which supports the idea that posture should be managed before discomfort builds up over time¹⁰.

The vision from a greater lens for Phase 2 is to make the chair understand posture as a **behavior**, not just a position. Phase 1 classified what the user’s posture looked like at a specific moment. That is useful, but posture is not one fixed moment. It changes slowly throughout a session. Phase 2 should look at trends, habits, and posture drift. Human mannerisms and behaviour are something that was barely mentioned in the first study, but is something that is critical for personalization. For example, if a user usually starts slouching after a certain amount of time, the chair should learn that pattern. If pressure gradually shifts toward one side, the system should recognize that early. In simpler words, the chair should stop thinking only in snapshots and start thinking in patterns. From a more holistic approach, the chair should use its own brain to “predict” what will happen before the user even does it.

Another major part of the Phase 2 vision is making the AI aspect more advanced, but also more practical. The goal is not to add AI just to make the product sound more impressive. The goal is to use AI in a way that actually improves the user’s experience and long-term sitting behavior. Random Forest was a strong starting foundation for Phase 1 because it worked well with structured pressure-sensor data and helped classify posture based on different sensor

readings. However, Phase 2 should use that as the baseline, not the final destination. The next version should track model confidence, posture duration, pressure changes over time, and user-specific baselines. Random Forest is known for combining multiple decision trees and averaging their outputs to improve prediction accuracy and help control overfitting, which makes it a strong MVP model, but the future of S.E.A.T. should move beyond basic classification and more into prediction and personalization.

The transformation from essentially furniture to a smart AI-leveraged product is really what makes Phase 2 important. A normal chair is passive and waits for the user to sit correctly. A smart ergonomic chair gives support, but it still depends on the user noticing their own posture and adjusting manually. S.E.A.T. has the chance to become something different because it can learn from the user and respond based on real behavior. In Phase 2, the chair should understand each user's baseline, detect when that baseline starts changing, and then choose the right type of intervention. Phase 1 defined intervention to be strictly physical, but that intervention becoming non-physical is exactly the shift from Phase 1 to Phase 2. That intervention does not always need to be a physical adjustment. Sometimes it could be a soft reminder, a small coaching message, or a suggestion to take a break. This makes the product feel more human because it is not forcing the user to adjust every time something changes. It is guiding them in a smarter way

Phase 2 also needs to make the product more realistic from a safety and trust compliance perspective. Since the chair has the ability to physically move while someone is sitting on it, the system cannot act carelessly and has no room for error(to prevent a lawsuit). The chair should not move every time the model predicts a posture shift, because people naturally shift throughout the day. Some movement is normal, and the product needs to understand the difference between normal movement and a sustained, unhealthy posture pattern. This is where confidence-based correction and false-correction prevention become essential. If the system is highly confident and the posture issue has lasted for a meaningful amount of time, then a physical correction may make sense. If the system is uncertain, then it should start with a softer intervention. This keeps the chair supportive and even gains the trust of a user instead of being annoying or unsafe. The product being annoying and unreliable is the biggest risk when looking at things from that regard.

The societal impact of Phase 2 is also a major part of the vision. This product is not just about creating a cool engineering system. It is about leveraging modern-day technology to address a problem that affects students, workers, remote employees, and people who sit for long periods every day. Multiple articles explain that ergonomics focuses on designing work environments around people to prevent discomfort and injury, which connects directly to what S.E.A.T. is trying to do. The stronger version of this product would not only help users sit better in the moment but also make them more aware of their habits over time. Ultimately, while the goal of this product is to gain monetary value, a huge win would be if a user fixes their posture for the better in the long run. Overall, Phase 2 should prove that S.E.A.T. is not just a chair with technology added to it, but a product that uses AI, hardware, and software in a meaningful way to improve how people sit, work, study, and protect their long-term health.




Image created using nano banana - Gemini

Image depicts an example of gradual posture decline and the intervention system of S.E.A.T

Phase 2: Enhancements

Phase 2: Technical, AI, and Data enhancements

General gist

- **Technical goal:** Make S.E.A.T. more reliable, standalone, smooth, and product-ready
- **AI goal:** Move from basic posture classification to predictive, personalized posture intelligence
- **Data goal:** Improve validation, reduce overfitting, collect better user data, and build a system that learns over time

This section is arguably the most important part of Phase 2 and the meat of the report because this is where S.E.A.T. moves from “the MVP that works” to “this product can actually become operable and has modern-day usage.” Phase 1 did the foundational work by proving that pressure sensors, machine learning, motors, and an application could work together. That is already a strong base. However, phase 2 needs to be 10-fold more extensive and detailed. Phase 2 should not just be about adding better parts or making the app look more user friendly. It should be all about improving and enhancing the full intelligence of the system: how the chair collects data, how the AI makes decisions, how the motors respond, how the user is protected, and how the product keeps improving over time.

The first major technical enhancement should be moving from basic posture classification to **predictive posture intelligence**. In Phase 1, the model focused mainly on the snapshot of the user’s pressure sensor readings and classified the posture at that moment. That is useful, but posture is not just a single, sudden moment. It is a behavior and mannerism that changes slowly. A user may sit upright at first, then slowly lean forward, lose back support, or shift more weight to one side once they start to feel fatigued. Phase 2 should track this posture drift over time. Instead of only saying, “the user is slouching,” the system should begin to say, “the user’s posture is starting to decline, and intervention may be needed soon.” This is a major shift because it changes S.E.A.T. from a reactive correction product into a proactive preventive support system.

To make this possible, the data going into the model needs to be improved massively. Phase 1 used pressure sensor values, which were the right starting point. But Phase 2 should go beyond raw sensor values and create stronger engineered features. For example, the system should track left-right pressure ratio, front-back pressure ratio, backrest-to-seat pressure ratio, pressure changes over multiple seconds and timeframes, posture duration, movement patterns, and how often the user transitions between posture states¹⁰. These features would help the model understand not only where pressure is located, but how the user’s sitting behavior is changing. That is where the shift from phase 1 to phase 2 occurs. A model that acts on posture change to a model that understands posture change.

Random Forest played a vital role as the machine learning model, but now it will strictly become the baseline, instead of the driver of this product. It was a strong MVP choice because it works well with structured sensor data, is easier to explain, and can classify posture based on

the 10 pressure sensor readings. With that in mind, Phase 2 should compare it against stronger tabular models, such as **Gradient Boosting**, because boosting models learn differently. Random Forest builds many decision trees independently and averages their decisions. Gradient Boosting builds trees one after another, where each new tree focuses on fixing the mistakes of the previous trees. To put it in simpler terms, Random Forest is like asking a group of people to vote at the same time, while Gradient Boosting is like having each person learn from the previous person's mistake before giving the next answer. For S.E.A.T., this could especially help when the model struggles with similar posture patterns, such as slouching versus leaning forward, or upright with back support versus upright without back support. A boosting model may be better at learning these smaller differences because it keeps improving on the errors the earlier trees made.

Phase 2 should also eventually test lightweight **time-series models**, but only after more long-session data is collected. The reason is that posture is not just a single sensor snapshot. It is a pattern that changes over time. Random Forest and Gradient Boosting can classify the current pressure pattern, but they do not naturally understand how the user's posture has been changing over different time frames, whether in minutes or seconds, unless we manually create time-based features. A time-series model is different because it is designed to look at sequences. With S.E.A.T. in mind, it could analyze how the pressure readings change across time and identify whether the user is slowly drifting toward slouching, leaning to one side, or losing back support. This would be extremely useful because the goal of Phase 2 is not just to detect bad posture after it happens, but to predict posture decline before it becomes a bigger issue and mitigate it in its tracks of occurring.

When looking at AI improvements, **confidence-based correction logic** definitely stands out. In Phase 1, once the model predicted a posture category, the system could use that prediction to trigger a correction. The issue is that not every prediction should automatically lead to movement. If the model is only slightly confident, the chair should not physically move under the user because that can become extremely annoying and even unsafe in some scenarios. Phase 2 should introduce confidence thresholds so the chair has different levels of response depending on how confident the model is. For example, if the model is highly confident, maybe around 85–90%, and the posture issue has lasted long enough, then the system can allow a physical correction. If the confidence is in the middle range, maybe around 60–85%, the app can show a reminder, or there could be a very slight physical correction, indicating posture fixture. If confidence is low, the system should keep observing and collect more sensor readings before doing anything. This makes the chair safer and smarter, and more importantly, gains the trust of the user to make good, accurate decisions.

To make this work, there needs to be accurate output labels. Instead of only saying “slouching” or “leaning left,” the system should also return a confidence score for each prediction. For example, the model might say: **slouching: 88%, leaning forward: 7%, upright: 5%**. In that case, the system can trust the prediction more. But if the model says **slouching: 52% and leaning forward: 45%**, then the chair should recognize that the decision is uncertain and avoid moving immediately. This would require the software layer to include a decision engine that sits between the AI model and the motor system. That decision engine would ask

three simple questions before allowing movement: **How confident is the model? How long has the posture issue lasted? Is the pressure trend getting worse or staying the same?** Only when all three conditions are met should the motor system physically adjust. This is a huge Phase 2 improvement because it makes the chair act more like an intelligent assistant instead of a machine that reacts too quickly.

Another important enhancement is **false-correction prevention**. Everyone moves around when sitting. People can lean, stretch, reposition themselves, cross their legs, adjust their back, and move around throughout the day. That does not automatically mean they have bad posture. A person leaning for two seconds is normal. A person leaning for twenty minutes may be an issue. Phase 2 should include a rule that combines model confidence, posture duration, and pressure trend before triggering physical movement. For example, if the model detects leaning left for only three seconds, the system should ignore it because it is probably normal movement. If the same posture continues for 45 seconds or more and the pressure pattern keeps drifting in the same direction, that is when the system can intervene for a posture issue. This helps the chair understand the difference between temporary movement and sustained poor posture.

To implement this, the chair would need a short-term memory layer that stores recent sensor readings and predictions over time. Instead of only looking at this from a holistic perspective, the system would look at a rolling window, within the range of 10-200 seconds. It would then try to classify whether this is a natural human movement or a pattern that keeps repeating towards bad posture. If the user briefly leans forward and then returns upright, the system should do nothing. If the user leans forward and stays there, then the system can respond gradually. This would stop the chair from overreacting to normal human movement and would also reduce unnecessary motor use. Honestly, this is one of the most important improvements because a product like this can fail quickly if users feel like it is constantly correcting them for small movements, which could become very annoying and turn users away. The best version of S.E.A.T. should feel supportive and calm, not jumpy or controlling.

Phase 2 should also introduce **closed-loop motor feedback**. In Phase 1, the system could predict posture and send motor commands, which proved the proof of concept. However, a stronger system should confirm whether the correction actually worked after the motor moves. Closed-loop control would take this a step further, and it means the system uses feedback to compare the actual result against the intended result, instead of sending a command and assuming everything worked. In simpler terms, it is like steering a car. You do not turn the wheel once and hope the car stays perfectly in the lane. You keep looking at the road and making small corrections based on what actually happens. Closed-loop systems use feedback from the output to adjust future action, instead of taking action blindly¹³.

In terms of S.E.A.T., the goal is to observe poor posture, act on it, and then understand if the user actually made any changes. If the pressure pattern moves closer to an upright or balanced position, the system stops. If the posture does not improve, the chair can either make a smaller secondary correction or switch to a softer intervention like an app reminder. This is

meaningful because it makes every movement measurable. The chair is no longer just saying, “I moved, so I assume I fixed it.” It is saying, “I moved, checked the result, and confirmed whether the correction helped.” This makes the chair smarter and more efficient, all while not being too pushy.

The motor system itself should also become more **comfort-aware**. This means the chair should not only know whether to move, but also how gently, how far, and how often it should move. Since the product physically moves while someone is sitting on it, the motion cannot feel sudden or aggressive. This could raise major concerns in terms of physical safety and could raise compliance issues. Phase 2 should include **micro-adjustments**, which means very small movements instead of big shifts. It should also include **slow ramp-up and ramp-down motor control**, meaning the motor starts moving slowly, gradually reaches the needed movement speed, and then slows down again before stopping to make the process feel smooth and slow. In simpler terms, it should feel like a smooth car braking gently, not a car slamming on the brakes.

This is important because the product is meant to improve comfort, not create new discomfort. In the case of S.E.A.T., that means the chair should adapt around the user instead of forcing the user into sudden or aggressive movements. This would lead back to the idea of personalization. Phase 2 should also include different features such as manual takeover and a cooldown period, which would again give the power in the hands of the user. Cooldown periods would stop repeated corrections from happening too quickly. A manual override would allow the user to stop or reject an adjustment if it feels uncomfortable. In simpler terms, the chair should feel supportive, not controlling. The user should feel like the chair is helping them with their posture in a smooth and helpful way, not take over and make the user feel any discomfort.

This would be the shift in making this feel like a purchase-worthy product and not just a working MVP. Another important technical improvement is edge-based processing. In simple terms, edge processing means running the AI model directly on the chair instead of depending on a connected laptop or outside computer. If the model runs directly on the chair, the system can respond faster because it does not have to send data back and forth to a connected laptop. This reduces latency, which basically means reducing the delay between the user’s posture change and the chair responding. It also improves portability because the chair can work in a classroom, office, library, or home without needing a full external computer setup. Google’s LiteRT, built on TensorFlow Lite, is designed for high-performance on-device machine learning and edge deployment, which shows that running smaller optimized models locally is a realistic direction for a product like this¹⁴.

The data strategy also needs to become much stronger in Phase 2. A high accuracy score is useful, but that does not mean that the product will work completely accurately on a test case it has never seen before. This is one of the biggest risks with a product like S.E.A.T. because different users have different heights, weights, body shapes, and sitting habits. The model needs to be tested in a way that proves it can generalize beyond the original training

data. For that, the training data needs to be very diverse, across thousands of different people, all with different factors mentioned above.


One way to achieve this goal is through leave-one-user-out validation. This means the model is trained on every user except one, and then tested on the user it has never seen before. In simpler terms, it is like studying with one group of people and then taking the test on a completely new person. If the model still performs well, that gives much stronger confidence that S.E.A.T. can work in the real world. Phase 2 should also collect longer sitting-session data instead of only short posture samples. The system should track how posture changes after hours instead of minutes because the intended ICP is aiming to target people who sit for long hours. This would help the chair understand fatigue, posture drift, and long-term behavior more accurately.

Another realistic enhancement would be building a personalized user posture profile. This does not need to be overly complicated. To put this in simple terms, this is understanding and personalizing the patterns of different users and understanding how each person's posture gradually worsens, rather than assuming. For example, one user may start slouching after 25 minutes, while another may lean right during long coding sessions. In simple terms, the chair stops treating every user the same and starts learning the individual sitting habits of each person. This builds directly on the original calibration idea, but takes it much further and gains the idea of trust.

Finally, Phase 2 should include a basic model governance process. This sounds technical, but the idea is simple: if the system is learning, retraining, or personalizing itself, the team needs to know which model version is being used, when it was trained, and whether it actually improved the process. This matters because S.E.A.T. is not just making a digital suggestion. It is making decisions that could physically affect the user. NIST's AI Risk Management Framework explains that trustworthy AI systems should be valid and reliable, safe, accountable, transparent, explainable, privacy-enhanced, and fair. For S.E.A.T., this means the AI cannot just be accurate once in a demo. It needs to remain controlled and trustworthy over time¹⁵.

Overall, the technical, AI, and data enhancements in Phase 2 should push S.E.A.T. beyond the original MVP, all while making the roadmap realistic, achievable, and, to put it bluntly, a product that has real-world value. Phase 2 should prove that the system can understand posture over time, act only when it is confident, move safely, work without a laptop, and improve through better data. That is how S.E.A.T. becomes more than a chair with pressure sensors. It becomes a smarter ergonomic system that feels useful, valuable, and most importantly, trustworthy for everyday users.

(a) Random forests



(b) Gradient boosting





Image showing the difference between random forest vs gradient boosting models



The image shows the workflow of closed-loop motor feedback

User Experience

User Experience, Safety, Privacy, and Human Factors

General gist

- **User Experience:** Make the product user-friendly
- **Safety:** Smooth-flowing physical movements
- **Privacy:** Use posture data responsibly without turning the chair into a surveillance product

A huge part of Phase 2 is making sure S.E.A.T. is not just technically impressive but also comfortable and trustworthy for the actual user. Phase 1 proved the concept of a foundational MVP working; however, a product like this cannot be judged only by whether it works in a demo. It also has to be judged by how it feels when a person uses it for a long period of time. The chair fails to do its job on a psychological level when it moves too aggressively or constantly, or makes the user feel like the purchase was not worth it.

The product should give complete ownership to the user of when the chair is actually in a position of moving them physically. S.E.A.T. should not have only one correction mode. Different users interact differently with the chair. While some might want soft correction, some might want app reminders, while some may even want hard physical correction. Phase 2 should include multiple modes, such as Passive Mode, Coach Mode, Assist Mode, and Auto Mode. Passive Mode would only track posture without correcting. Coach Mode would give reminders and feedback. Assist Mode could ask before physically adjusting. Auto Mode would allow automatic correction once the user trusts the system. To put it in simpler terms, the user should decide how involved the chair is in their posture correction process.

A chair that can move you against your will can raise a lot of compliance issues, and that is why building trust is essential. Neutral body positioning helps reduce stress and strain on different muscle groups, tendons, and the skeletal system, but reaching that position should still feel natural and comfortable for the user. The goal should not be to force the user into one “perfect” posture every second. Most people move, shift, and adjust throughout the day at different times of the day. S.E.A.T. should support better posture without making the user feel trapped or controlled.

The idea introduced plays itself perfectly for a hierarchy system. Instead of jumping straight to motor movement, the chair should start with the least invasive option. For example, when the system notices a small posture issue, it could first show a visual cue on the app. If the issue continues, it could send a reminder. If the user still does not adjust, the chair could make very small micro-adjustments to gain the attention of the user. Full physical correction should be used only when the system is confident, and the posture issue has lasted long enough. This makes the product feel calmer and more human instead of a machine. It also reduces unnecessary motor use and makes the chair less likely to annoy the user.

Safety also needs to become a serious Phase 2 focus. Because S.E.A.T. has the ability to physically move while someone is sitting on it, the system needs clear safety rules. These should include manual override, maximum movement limits, motor cooldown periods, and a safe default position. A manual override lets the user stop or reject a correction if it feels uncomfortable. This lets the user feel in control at all times and mitigates safety issues. Maximum movement limits prevent the chair from moving too far. Cooldown periods stop the chair from repeatedly correcting the user too quickly. A safe default position means that if the system becomes confused, loses sensor confidence, or detects an issue, it stops moving and returns to a neutral state.

Human factors also matter because people are not machines. Two users can sit differently and still both be comfortable. A taller user, a shorter user, an athlete, a student, or even an office worker may all have different posture patterns and different comfort preferences. Phase 2 should avoid treating every user the same. The chair should learn each user's baseline and allow them to adjust sensitivity in the manner they want to. For example, one person may want stronger correction because they are trying to fix a long-term habit. Another person may only want light reminders because they find movement distracting. Someone young can feel comfortable with hard correction, while someone older might need softer correction. This level of control would make the chair feel more personal and less generic.


When commercial institutions are looking into S.E.A.T, there should be a clear separation between institutional and personal insights. For example, an employer or university should not see that one specific student or employee slouched for a certain number of hours. Instead, the institution should only be able to access generic data, such as average chair usage, general posture improvements and trends, or maintenance needs. NIST's Privacy Framework explains that privacy risk management helps organizations build innovative products and services while protecting individual privacy¹⁶. That idea fits S.E.A.T. well because the product needs data to improve, but it also needs strong boundaries so users feel safe using it and not feel like they are being seen or judged by other people.

Overall, the Phase 2 user experience should make S.E.A.T. feel more supportive, personable, safe, and most importantly, trustworthy. The chair should not just be smarter from a technical perspective but also feel like a privilege to use. That means giving users control, using softer interventions before physical correction, limiting risky movement, personalizing comfort settings, and protecting posture data. Phase 1 proved that the chair can work. Phase 2 needs to prove that people use this product in their everyday lives.

7 ELEMENTS OF PRIVACY BY DESIGN

A holistic framework for integrating data protection into products, services, and system designs by default.

This strategic approach to managing risk and creating a better experience for users, is a requirement under numerous Privacy regulations.



This image depicts the privacy framework

The Business Strategy

Business Model, ICP, and Go-To-Market Strategy

General gist

- **The Business Model:** From selling a “smart chair” to selling a “hardware + software” ergonomic platform
- **ICP:** universities, computer labs, libraries, and corporate office environments
- **GTM-Strategy:** Launch through pilots, prove measurable posture improvement, then expand into larger institutional deployments

This, alongside the technical enhancements are arguably the most important factor of phase 2 because creating a working MVP is impressive, but creating a commercial phenomenon takes PM-level thinking. Phase 1 proved that the product can work as a technical MVP, but a working MVP is not automatically a sellable product. For S.E.A.T. to become more than a capstone prototype, the business side has to be just as strong as the engineering side. The product needs a clear customer, a clear reason to buy, a clear pricing model, and a realistic path to adoption for acquisition. In simpler terms, Phase 2 needs to answer one very simple yet major question: **who would actually pay for this, and why would they?**

The first major mistake would be trying to sell S.E.A.T. to everyone. Technically, almost anyone who sits for long periods could benefit from it, including different consumer targets such as students, office workers, remote employees, gamers, and patients in rehabilitation. However, when you look at it from this perspective, this creates a wide variety of range of people to sell to. A strong product strategy always has a sharp and clearly defined ICP for the marketing aspect. The best initial starting point for the ICP for S.E.A.T. should be environments where people sit for extremely long periods, where posture problems are common, and where one buyer can purchase multiple units. Putting this idea in mind, these points toward universities, computer labs, libraries, corporate offices, call centers, and maybe the potential idea of rehabilitation clinics with Phase 3 in mind later.

The strongest initial market fit would likely be universities and computer labs to start off. Additionally, this makes sense because students sit for hours while studying, coding, doing labs, or working on assignments. Universities also have the idea of centralized purchasing, which can mean the product is bought in bulk, wellness initiatives for their staff and students, innovation programs, and facilities budgets. Additionally, since S.E.A.T. came from a university engineering environment, the first proof point should naturally come from that same setting. A university pilot would be easier to justify than a full commercial launch because it can be framed in multiple ways. A wellness initiative and a student-led innovation project that is being tested for proof of concept.

The second strongest market fit would probably be the corporate setting and call centres. These environments usually have large groups of people sitting for long periods, and companies need to care about productivity, comfort, workplace experience, and employee wellness. The National Safety Council also reported around 937,620 musculoskeletal disorder

DART cases in 2023–2024¹⁷. This shows the growing demand for a product like S.E.A.T., which can be sold in millions with the right idea. It can be positioned as a workplace wellness and ergonomic risk-reduction tool, which can give an incentive for these employees to work for longer hours without having to worry about developing chronic back pain in their later years.

From a business model standpoint, Phase 2 should move away from the idea of a furniture or hardware sale. Looking at this from a perspective of smart furniture will hold the product back from the astronomical heights this product can go to if marketing is done properly. Selling the chair alone can generate revenue, but it also limits the product's long-term value. The stronger model is positioning this as a **hardware-plus-software platform**. The hardware is the smart chair itself, but the software creates ongoing value through posture analytics, user profiles, calibration, maintenance alerts, and institutional reporting. This is what turns S.E.A.T. from a product someone buys once into a platform that can keep improving over time, and keeps investors and institutions ready to buy this product in bulk.

The basic revenue model could have three different parts. First, there is the “hardware” chair purchase itself. Second, there is a “software” subscription for organizations that want different features like analytics, dashboards, calibration management, and reporting. Third, there can be a reward-based support or maintenance system for institutions that purchase multiple chairs. This will help early sales and get the product in hot demand. For example, if a university was to buy 10-20 chairs in a pilot testing environment, they would receive a dashboard with all analytics, a setup support system, and reports of posture improvement. That is much more convincing than simply asking someone to buy a chair because it is “smart.”


The go-to-market strategy should definitely look into piloting instead of explicitly selling the hardware product. A 60–90 day pilot would allow S.E.A.T. to prove whether users actually like the chair, whether the chair improves posture behavior, whether the hardware can survive daily use, and whether the analytics are useful. During the pilot stage, the key metrics should include average daily usage, calibration completion rate, posture score improvement, slouching time reduction, user comfort rating, manual override rate, and system reliability. Over time, these metrics can also be used in marketing to show the general audience the benefits of a chair like this. These are the numbers that would help turn S.E.A.T. from an interesting prototype into a product with evidence behind it.

The first pilot should be small and controlled. A realistic starting point would be 5 to 10 chairs in a Rutgers computer lab, engineering lab, library area, or a common study space. The team could collect baseline posture behavior for the first week, then compare it with the results of later weeks after reminders and corrections are turned on. The final pilot report could show many different metrics such as posture fixtured rate, the average usage of the chair, and whether the general audience is actually satisfied with the product or not. This would give the team valuable data and a study that S.E.A.T. can then use to expand to more university spaces, then to nearby universities, and then eventually to corporate office environments.

The buyer-facing and the user-facing propositions are 2 completely different stories. For the user, the message is simple: **sit with better support and understand your posture habits so you do not have to deal with health issues down the line.** For the buyer, the message is different: **improve ergonomic wellness, collect privacy-safe aggregate insights, support healthier sitting environments, and ultimately increase the productivity of your employees.** This ties back to the business case that posture and workstation design are real concerns, and not just personal preferences.

Privacy is a huge part of the GTM strategy, especially when commercially aiming for workplaces and schools. If an employer or university buys the product, users should not feel like they are being watched. The institutional dashboard should only show anonymized aggregate data, instead of individual sitting behavior. For example, the dashboard could show average posture improvement across a lab, chair usage rate, calibration completion rate, maintenance alerts, and even satisfaction levels of the chair. It should not show that one specific student or employee slouched for two hours or anything related to the personalized bad posture of an individual. This privacy-first approach can become a selling point because it gives institutions useful insights without creating a surveillance problem.

Overall, the Phase 2 business strategy should focus on turning S.E.A.T. from a working prototype into a product with a real adoption path. Ultimately, if the product does not have a strong business backing and a standardized plan, all the hard work and engineering would go to waste. The goal should not be to immediately sell to every possible customer. The smarter move is to start with a focused ICP, run pilots, prove measurable value, and then expand once initial success is found. If S.E.A.T. can show that it improves posture behavior, feels comfortable to use, protects privacy, and provides useful institutional insights, then it becomes much more than a smart chair. It becomes a sellable ergonomic platform with real demand in a rapidly growing niche market.



This image depicts the importance of finding the right ICP and GTM strategy

Manufacturing + Scalability

Manufacturing, Cost, and Scalability Plan

General gist

- **Manufacturing:** Prototype to fast production
- **Cost:** Cost analysis
- **Scalability:** assemble, repair, deploy, and maintain across many users

Another major part of Phase 2 is making S.E.A.T. more realistic from a manufacturing and scalability standpoint. A working MVP is much different than a product that can be assembled in millions, shipped, and delivered to users around the world. This is where Phase 2 needs to mature the product. The goal is simple. It should aim to make the chair not only smarter but more accessible and reliable in a real world setting like office rooms and labs where hours are spent working.

The first manufacturing enhancement that stands out is the transition from an MVP level prototype to the transition of a repeatable and consistent system that is constantly iterated, assembled and shipped. It made sense in phase 1 that the team used available parts, exposed wiring, external hardware, and a more hands-on approach because the goal was to prove the concept. However, if S.E.A.T. is going to scale, the design needs to become more standardized and whole. This would mean cleaner wiring, stronger mounting, protected electronics, and parts that can be assembled the same way every time. NIST explains that moving a new product from idea to production involves prototyping, testing, selecting the right materials, and validating how well the product satisfies customer needs¹⁸. This connects directly to S.E.A.T. because the next version needs to move from a working build into something that can be manufactured consistently and frequently if the goal of the product is to see upscaled commercial success.

Cost analysis needs to be heavily stressed for the launch of a product. As of right now, the MVP's cost analysis is a little weak and unorganized. A real product cost includes more than sensors, motors, foam, Arduino, and mechanical parts. It also includes assembly labor, packaging, shipping, quality testing, software hosting, repairs, warranty support, and replacement parts. This is important because a product can look affordable at the prototype level but become much more expensive once it is manufactured and supported at large scale. Phase 2 should create a more complete cost model that separates prototype cost, production cost, and total product lifecycle cost.

One major cost and scalability improvement would be moving toward a custom PCB. A PCB is a board that holds and connects electronic components more cleanly and reliably in a more organized version. For S.E.A.T., this would help with the organization aspect by a lot. Connecting different sensors, motor drivers and all sorts of messy wiring would get much easier. What this would also do then is make the chair way easier to assemble and make repairs 10-fold easier. Additionally, the human error of wiring issues would also in turn be mitigated and the reliability aspect would get much better. Siemens explains that designing for


manufacturability in PCB design can streamline production, reduce errors, and save time and money¹⁹.

The whole computing setup also needs some work to become more product-ready. A real S.E.A.T. chair should not need a laptop connected to it to run properly. Phase 2 should move toward onboard computing, where a small embedded device runs the posture model locally. Raspberry Pi's Compute Module products are designed for embedded and industrial solutions, with the Compute Module 4 starting as low as \$41.25, which shows that onboard computing is actually way more efficient and can be realistic from a cost standpoint. For S.E.A.T., this would improve portability and make the chair feel like a real product instead of a lab setup or a test design²⁰.

Scalability does not only mean growth but also the repeated survival of the chair after thousands of testing logs. Phase 2 should include durability testing for sensors, motors, gears, wiring, cushions, and all the hardware component parts. This is especially important because the chair physically supports body weight and makes mechanical adjustments. The team should test how the chair performs after hundreds or thousands of sitting sessions, repeated motor movements, and different user weights. If the product is eventually placed in a computer lab or office, it will be used by many different types of people, not just one controlled tester. That means durability and consistency become just as important as model accuracy for the growth of a product like this.

Overall, the Phase 2 manufacturing plan should focus on making S.E.A.T. more repeatable, sustainable, realistic, and reliable. Phase 2 needs to prove that the chair can be built again and again without losing quality. That means cleaner electronics, a better cost model, modular repairable parts, onboard processing, stronger mechanical design, and real durability testing. In simple terms, the next step is moving from "we made one working chair" to "we can produce a version of this chair that people could actually use every day."

How to scale your digital product?



 **SIMFORM**

Image depicting a basic scaling idea for a product

Roadmap & The future

Roadmap & The future

General gist

- **Roadmap:** Realistic building in an organized fashion
- **Validation:** Proof of concept
- **The future:** Potential ideas if Phase 2 finds success

The roadmap for Phase 2 is a simple one. The goal is to turn an MVP into a scalable and reliable product bought and subscribed to by users around the world. The biggest mistake would be trying to build every idea at once. Phase 2 needs to be ambitious, but it also has to be realistic. The more efficient approach is to break this down in multiple stages so everything can be organized and approached head on. First stabilize the MVP, then improve the intelligence layer, then make the chair more product-ready, and finally test it in a real environment.

The first stage of this should be **MVP stabilization**. In the context of S.E.A.T., this would mean cleaning up the current system, solidifying it by testing the product over thousands of iterations and lastly making sure the existing features work consistently before adding too many new ones. The team should prioritize the data pipeline initially, reduce bugs in the app, organize the calibration process, and make sure sensor readings are stable so the chair can work at maximum accuracy. This is also the stage that should include basic safety rules like manual override, maximum movement limits, and a safe default position so the user does not feel overwhelmed by the chair.

The second stage then should focus on the **AI and correction logic**. This is where Phase 2 starts becoming more advanced and adds the element of a “smart” product. The chair should move beyond basic posture classification and begin tracking posture drift, model confidence, and posture anomalies with their respective durations. Instead of reacting every time the model predicts a posture, the system should ask whether the prediction is confident, whether the posture has lasted long enough, and whether the user is naturally correcting themselves. This stage should include confidence-based correction, false-correction prevention, and a rolling time window that tracks posture behavior over several seconds or minutes and then gradually into hours. To explain it in simple terms, the chair should become calmer and more thoughtful before it decides to move.


The third stage should focus on **product readiness**. This is making this feel less like a lab tested MVP to more of a everyday product to be used. The system should move toward onboard or edge-based processing, where the AI model runs directly on the chair instead of depending on a connected laptop. Google’s LiteRT is designed for high-performance on-device machine learning and edge deployment, which supports the idea that local AI processing is a realistic direction for a product like S.E.A.T. This would reduce delay, improve portability, and make the chair easier to use in classrooms, offices, libraries, or homes.

The fourth stage should be a **real-world pilot**. This is arguably the most important stage. Builders can engineer a perfectly functioning chair but if it has no real public demand, the product is useless. A realistic pilot could involve 5 to 10 chairs placed in a Rutgers computer lab, engineering lab, library, or study space for 60 to 90 days. During this pilot, the team should track average usage, calibration completion, posture improvement, false-correction rate, manual override frequency, user satisfaction, comfort rating, and system uptime. The goal is not just to prove that the technology works. The goal is to prove that people trust it, like it and would want to use it and integrate it in their everyday lives.

Risk mitigation should be built into every stage of the roadmap. The biggest risks are overfitting, false correction, user discomfort, hardware failure, privacy concerns, and weak adoption. Each risk should have a clear response. The issue of overfitting can be reduced through more diverse data and leave-one-user-out testing as explained in previous tabs. False corrections can be reduced through confidence thresholds and posture-duration rules. User discomfort can be reduced through micro-adjustments and manual override. Privacy concerns can be reduced by using anonymized aggregate data for institutions instead of individual posture tracking. Once every risk has a proposed solution, that is when the product can see serious success atleast from a technical and financial standpoint.

After Phase 2 proves reliability, Phase 3 can focus on bigger future opportunities. One possible Phase 3 idea is a broader workplace or campus ergonomic platform, where multiple chairs connect to an institutional dashboard showing aggregate posture trends, chair usage, and maintenance alerts. Another possibility is a sense of deeper personalization, where the chair learns to treat the user as a person and understands their sitting and posture mannerisms and behavioral movements. A future phase 3 version could also expand into the idea of rehabilitation or physical therapy settings, but this would need constant iteration since healthcare is a very volatile industry and leaves no room for errors. Healthcare use cases sound attractive, but they should not be the first step because they require more significantly more evidence, trust, and caution due to serious compliance issues.

Overall, the roadmap should keep Phase 2 focused and Phase 3 ambitious. Phase 2 should prove that S.E.A.T. can work safely, consistently, and most importantly, comfortably in real environments. Phase 3 can then explore larger-scale deployment, deeper personalization, ambitious and innovative ideas for all sorts of users and new markets. In simple terms, Phase 2 is about making the MVP into a product. Phase 3 is about making the product bigger.



The image depicts a typical product lifecycle dev

Conclusion

Conclusion

General gist

- **Takeaways:** Phase 1 vs Phase 2
- **The Ending Message:** The creation of a working chair into the next multi billion dollar product

Overall, S.E.A.T. should be viewed as a strong foundation with a much bigger and bright future ahead of it. Phase 1 proved the most important first step: the chair can sense posture, classify it through machine learning, respond physically through motorized adjustment, and show user feedback through an application. That is not a small accomplishment. This portrays a project that was built over multiple months of dedication and a working MVP with the integration of hardware, software and AI systems. However, the real opportunity is not stopping at the MVP. The real opportunity is taking what Phase 1 created and transforming it into a product that is smarter, safer, more useful, and realistic enough to be used by real people every day.

The biggest takeaway from this report is the difference between proving a concept and building a product. Phase 1 answered the question, “Can this chair work?” The answer is yes it can. Phase 2 now needs to answer a much more serious question: “Can this chair work well enough to be trusted, scaled, sold, and used consistently?” That is where the project becomes more exciting. A chair that works once in a demo is impressive, but a chair that can work in a computer lab, office, library, or workplace every day is in a completely different stratosphere. This is why Phase 2 needs to focus on predictive AI, safer correction logic, smoother motor movement, stronger validation, better personalization, clearer business strategy, and scalable manufacturing.

The reason S.E.A.T. matters is that the problem itself is real. The opportunity is huge because healthcare is a field that will always have demand, and this is directly conquering that field. With chronic back pain projected patients is only going up, this gives this product a great chance of growing in an industry where health is widely recognized to be very important. This gives S.E.A.T. a real purpose because it is trying to use AI, hardware, and software to help solve a problem that affects students, workers, remote employees, and people who sit for long periods every day.

In the end, Phase 1 built the working chair. Phase 2 is about building the product around the chair. That is the real difference. The final version of S.E.A.T. should not just move when posture is poor. It should understand when posture is starting to decline, know when not to move, protect the user’s comfort, respect their privacy, and give institutions a reason to invest in it. If Phase 2 is executed correctly, S.E.A.T. can move from being a successful capstone MVP into a serious product direction with real-world demand. The ending message is simple: Phase 1 proved the chair can work, but Phase 2 is what can turn a foundational MVP into something worth building, selling and scaling.

Citations & Refernces

<https://www.who.int/news-room/fact-sheets/detail/low-back-pain>¹

<https://www.fortunebusinessinsights.com/ergonomic-chair-market-112042>^{2, 8}

<https://www.health.harvard.edu/healthy-aging-and-longevity/in-a-slump-fix-your-posture>³

<https://www.geeksforgeeks.org/machine-learning/how-decision-tree-depth-impact-on-the-accuracy/>⁴

<https://journals.sagepub.com/doi/pdf/10.1518/107118188786762441>⁵

<https://arduinoexpert.com/seating-posture-correction-device-using-arduino>⁶

<https://www.amaehealth.com/blog/posture-and-mental-health>⁷

<https://pmc.ncbi.nlm.nih.gov/articles/PMC6166197>⁹

<https://www.osha.gov/etools/computer-workstations/positions>¹⁰

<https://www.cardinus.com/us/reactive-vs-proactive-ergonomics>¹¹

<https://yentzfamychiropractic.com/bad-posture-and-treatment>¹²

<https://www.a-m-c.com/closed-loop-control>¹³

<https://ai.google.dev/edge/litert>¹⁴

<https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>¹⁵

<https://www.nist.gov/privacy-framework>¹⁶

<https://injuryfacts.nsc.org/work/safety-topics/musculoskeletal-injuries>¹⁷

<https://www.nist.gov/mep/product-design-and-development>¹⁸

<https://blogs.sw.siemens.com/electronic-systems-design/2024/02/07/design-for-manufacturability-in-the-pcb-design-process>¹⁹

<https://www.raspberrypi.com/products/compute-module-4>²⁰

The images used for the Strategic Overview and Phase 1 recap were created using Gemini