Cognitive Comfort Index: Technical Report

By Alexander Wilkins

Introduction

As a computer science and finance double major at the University of Georgia, I spend a lot of my time doing homework. Naturally, most of the hours I spend doing this homework are spent in my room in my cheap apartment in downtown Athens. This room lacks any sort of natural light, has bright LEDs shining down from the ceiling, gets incredibly stuffy, has paper thin walls and if I am working late at night on any day but a Monday you can feel the room shake from the vibrations of the bar downstairs. This very clearly is not the optimal environment for getting work done and I am certain my productivity takes a hit from one or all of the factors listed above. In my struggles with this study space, I began to wonder what the optimal conditions for productivity actually are, and whether I would be able to quantify the actual impact to my productivity when I am working in my room. With this in mind, I set out to create the Cognitive Comfort Index to measure a room's potential for productivity and to find the best possible study spot. The idea was to use a system of sensors to take readings on different environmental factors within the indoor ecosystem and compare them to a benchmark. By looking at these comparisons, I would then be able to combine them into a single score, or index, to determine the overall impact of such factors on a user's cognitive ability and define a room's capacity for productivity. The factors I was interested in analyzing the effects on an individual were the room's lighting, humidity, temperature, air quality, and noise.

System Design

In this section I give a more detailed overview of the system design for my project, covering both the hardware and software components as well as the core scoring methodology that drives the index.

My original plan for the hardware portion of this project was to integrate each sensor individually and connect them into a Raspberry Pi based system. Luckily, my professor was aware of the Enviro+, a device by Pimoroni that is designed for environmental monitoring. This device succinctly integrated all of the sensors I needed for this project into one easy to connect piece of hardware. This device, integrated with my Raspberry Pi 4, became the foundational hardware for my project. After setting up my Pi with the enviro-plus environment, I realized that none of the sensors were working. After extensive testing, I finally determined that somehow the Enviro+ was completely broken. I was unfortunately forced to buy a new one to continue my project and in another disappointing event I was only able to get my hands on an Enviro Mini, which did not contain any sensors for measuring air quality. This meant that I would have to focus on the effects of lighting, humidity, temperature, and noise, although these proved to be more than enough.

With this revised hardware setup in place, I focused on the design of a software system capable of reliably collecting, processing, and displaying environmental data in real time. The codebase is organized into three primary modules: a sensors layer, a scoring layer, and a Flask web server. The sensors layer, contained in `read_sensors.py`, handles all direct communication with the Enviro Mini hardware using the `bme280` and `ltr559` libraries for temperature, humidity, and light, and the `sounddevice` library for noise. The scoring layer, contained in `scoring.py`, receives those readings, applies research-derived scoring functions to each factor, and combines them into the Cognitive Comfort Index. The Flask web server sits on top of these two layers, exposing API routes that trigger fresh sensor reads on demand and serving a live frontend dashboard to any device on the same network.

A key challenge in the sensor layer was the temperature reading. Because the Enviro Mini sits directly on top of the Raspberry Pi, the heat generated by the CPU inflates the raw temperature reading considerably. To correct for this, I implemented a compensation formula that reads the CPU temperature directly from the system at `/sys/class/thermal/thermal_zone0/temp` and uses it to subtract the inferred CPU contribution from the raw BME280 reading. The specific formula applies a compensation factor of 1.5, which I tuned empirically by comparing the corrected output against a standard thermometer until the values were consistent.

Noise measurement presented a different kind of challenge. The Enviro Mini captures raw audio samples at 44,100 Hz over a 0.1 second window using the device's microphone. From these samples, the system computes the root mean square (RMS) amplitude and converts it to a decibel value using the standard formula, with a calibration offset of 70 dB applied to bring the output into a realistic absolute range. This offset is an approximation rather than a laboratory calibration and is one of the more variable aspects of the system, as discussed further in the Evaluation section.

Scoring Methodology

The scoring system is the intellectual core of the project. Each environmental factor is converted into a score from 0 to 100 using a function derived from published research, and the four factor scores are then combined into the overall Cognitive Comfort Index (CCI) through an equally weighted average.

Temperature is scored using a cubic polynomial derived from research by Lawrence Berkeley National Laboratory on the effect of indoor temperature on task performance. The polynomial maps Celsius values to a productivity percentage, which is then clamped to the 0 to 100 range.

Light is scored on a logarithmic scale relative to an optimal target of 500 lux, a threshold informed by research on illuminance and visual comfort. The logarithmic form was chosen because human perception of light intensity is itself roughly logarithmic, and because it allows the score to rise quickly at low lux values where the productivity impact is most severe and level off as conditions approach optimal.

Humidity is scored using a quadratic penalty function centered at 45% relative humidity, which represents the midpoint of the range most commonly cited in indoor air quality literature as conducive to both comfort and productivity. The penalty coefficient was tuned so that readings at

30% and 60% RH, representing mild deviation from optimal, yield scores near 95 out of 100. The research basis for this threshold draws from work published in the American Journal of Industrial Medicine and related occupational health literature.

Noise is scored using a piecewise linear model derived from Srinivasan et al. (2023) in npj Digital Medicine, which found that physiological wellbeing is optimal at approximately 50 dBA. Below 50 dBA, each additional 10 dB of quiet improves the score by 5.4 points, reflecting the finding that very quiet environments actually carry a small wellbeing penalty relative to a mild ambient background. Above 50 dBA, each additional 10 dB reduces the score by 1.9 points.

The four individual scores are combined into the CCI through a simple equally weighted average, with each factor contributing 25% to the total. This weighting decision was made deliberately. Given the mixed and often context-dependent nature of the research on the relative importance of each factor, equal weighting was the most defensible approach for a generalized productivity index. A more sophisticated version of this system could allow user-defined weights based on individual sensitivity.

Database and Leaderboard

The system persists sensor data using a SQLite relational database stored at `data/sensor_logs.db`. Each record in the `sensor_logs` table captures a UTC timestamp, a location name, the four raw sensor readings (temperature in Fahrenheit, humidity percentage, light in lux, and noise in decibels), and the five computed scores (one per factor plus the total CCI). The logger module averages readings collected over a five second interval before writing a record, which reduces the effect of momentary spikes in any single sensor. Users can save a named study spot through the frontend interface, which associates the current sensor snapshot and scores with a user-provided location name. These saved spots are displayed on a leaderboard within the dashboard, allowing comparison across different environments.

Evaluation

Throughout my journey with this project I ran into many setbacks that I had to think critically and creatively to push through. The breaking of the Enviro+ was a pretty low blow to begin the project with and while it was likely my fault, I am not exactly sure what caused the issue, but it was something I had to persevere through nonetheless. Despite the hardware issues, I am still very proud of my overall product and feel that my results validate the amount of work that I put in. I especially feel that the UI stands out and goes above and beyond to create an intuitive and functional environment for the user. It allows the user to simply and effectively interact with the data coming from the Enviro and even allows functionality across different environments through the SQLite database and the built in leaderboard.

In regards to the key functionality of the project, the actual productivity data and scores, the determination of these items was definitely the most difficult piece. For certain factors there was much more information available on their impacts on productivity than others. While some research I was able to find gave me actual formulas for their impact on productivity, others had very little quantitative information available. This made it difficult to determine the scoring functions or optimal conditions for some of the factors. In the case of temperature, I was fortunate to find a direct polynomial productivity model from Lawrence Berkeley National Laboratory that I could implement directly. Noise and humidity required more interpretive work, where I had to translate the qualitative findings of published studies into functional scoring curves. Light fell somewhere in between, where the 500 lux optimal threshold is well established but the exact shape of the productivity curve required judgment in selecting the logarithmic form.

In addition to the research challenge, the Enviro Mini trades simplicity and size for accuracy, and many of the readings, particularly noise, experienced a lot of variability in practice. The 0.1 second sampling window captures a very narrow slice of the acoustic environment and is sensitive to brief transient sounds that may not reflect the true ambient noise level. The calibration offset used to convert RMS amplitude to decibels is also an approximation, and I would not characterize the noise readings as precise in an absolute sense. With all of this being said, the system successfully achieved its core objective: creating a real-time environmental monitoring platform capable of measuring conditions, interpreting them through a research-driven scoring model, and presenting the results through a clean and responsive interface. While not perfect, this project serves as a minimal viable product and a baseline for my interest in this idea, and I look forward to continuing to build on it.