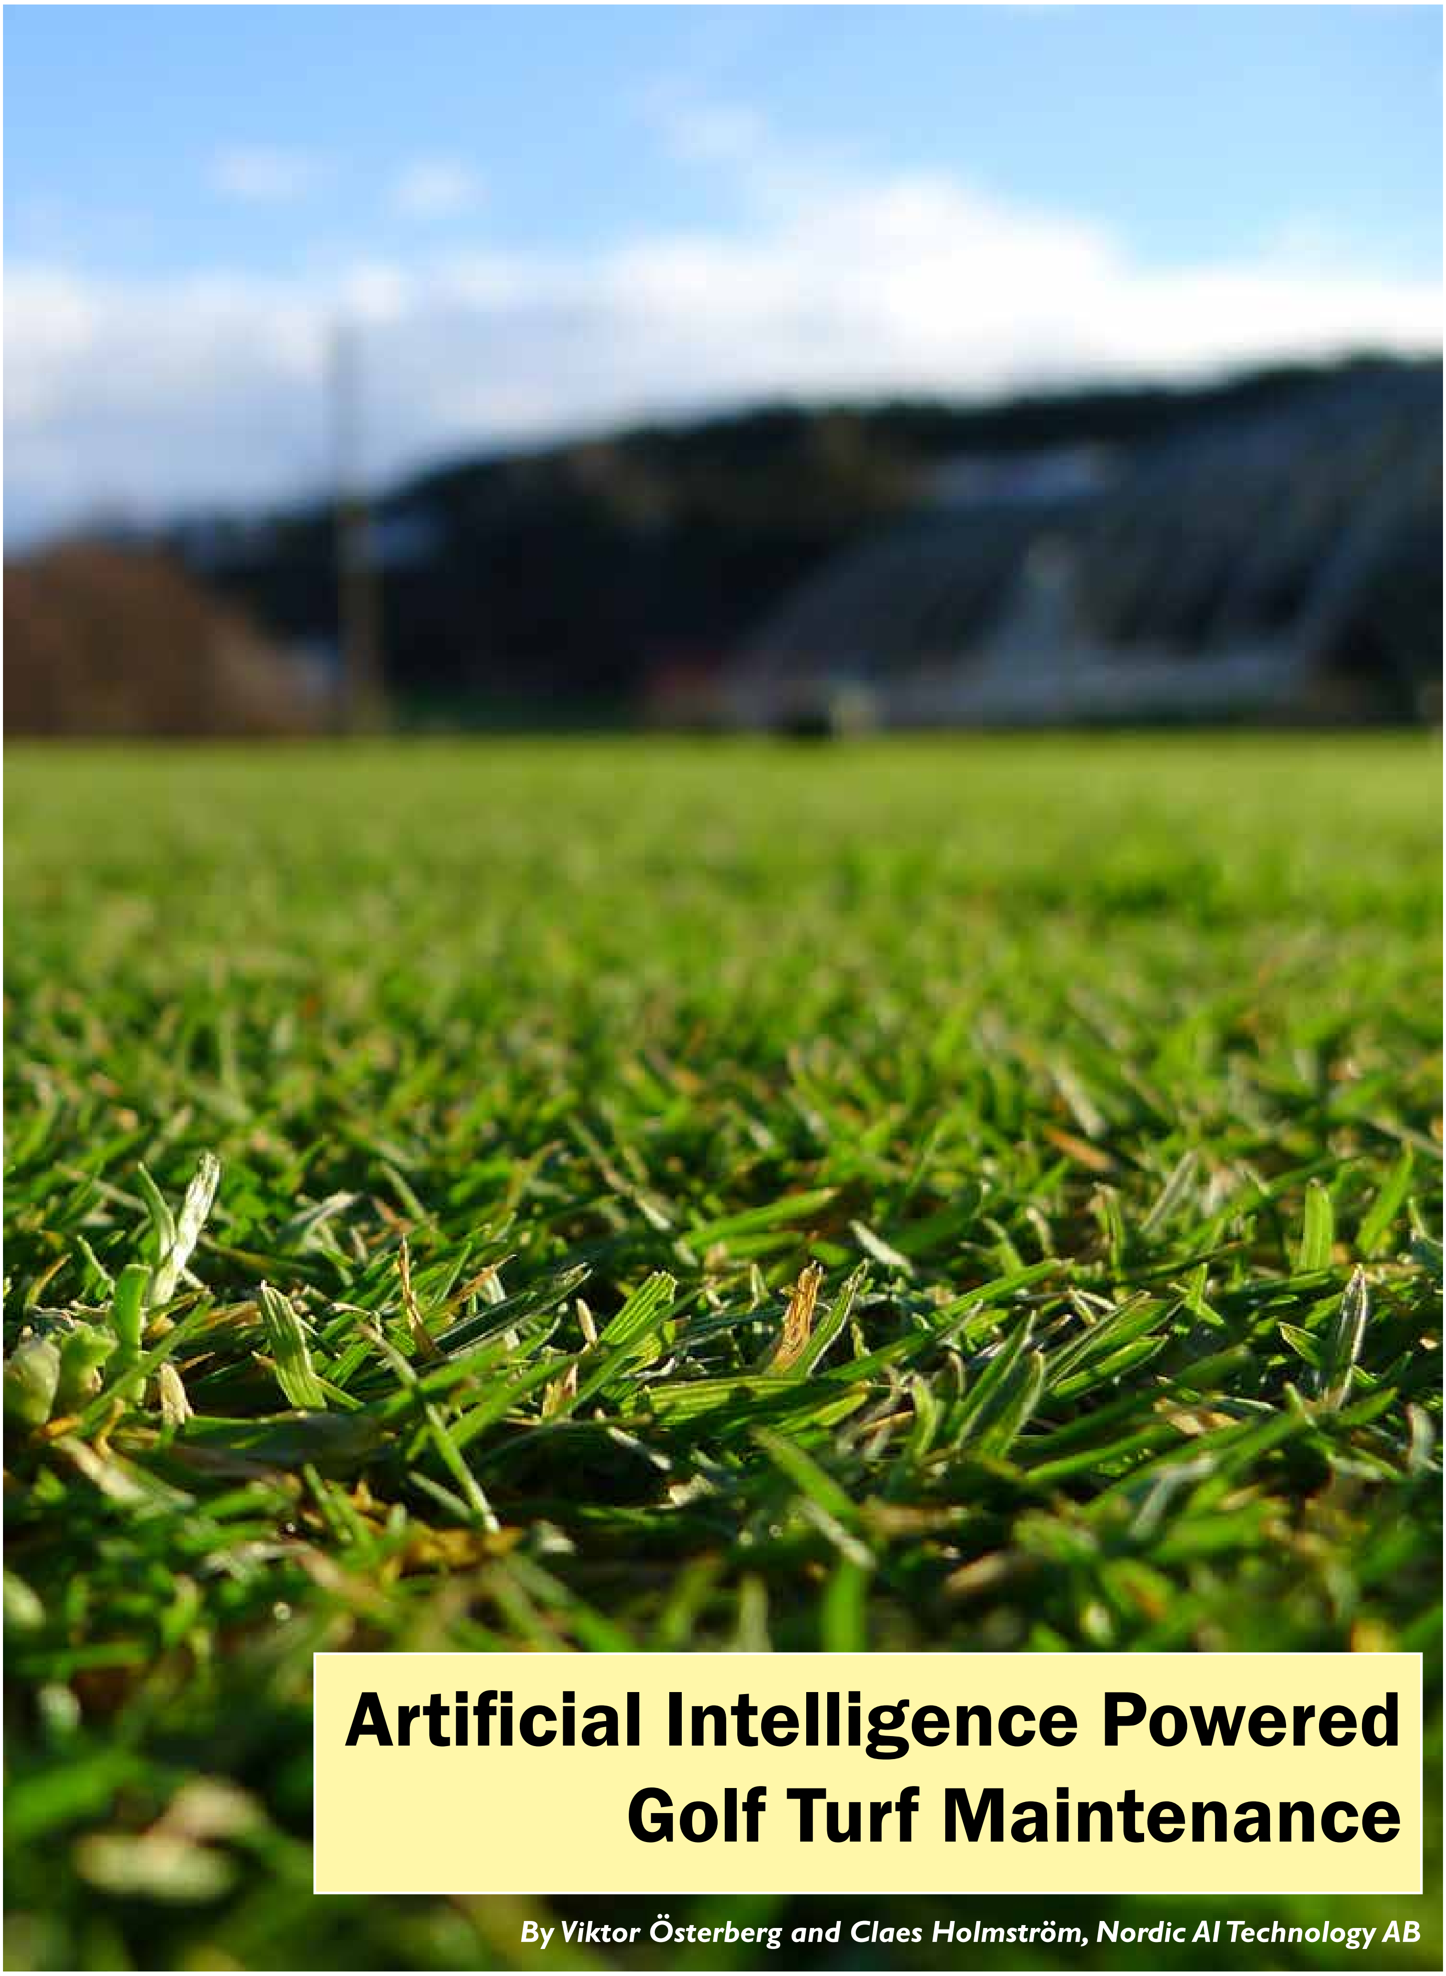

Artificial Intelligence Powered Golf Turf Maintenance

By Viktor Österberg and Claes Holmström, Nordic AI Technology AB

Popular Scientific Articles - STERF, May 2025

Artificial Intelligence Powered Golf Turf Maintenance

Aim of the Study

Golf courses are under pressure. Climate change, increasing operational costs, and tighter rules on pesticides are making it harder to maintain high-quality turf. At the same time, golfers expect fast, consistent greens, while the industry pushes for sustainability and reduced environmental impact. Traditionally, the job of managing turf has fallen to greenkeepers, who rely on years of experience and trial-and-error strategies. But is there a smarter way to do this?

The ADORE project set out to explore a modern solution to an age-old problem: Can artificial intelligence (AI) help golf course managers make better, more informed decisions about turf maintenance? More specifically, can AI work effectively using only the limited types of data that an average Scandinavian golf club might already have- things like basic weather data, maintenance logs, and occasional soil analysis?

The research team behind ADORE, which included experts from the Royal Institute of Technology KTH and Nordic AI Technology, wanted to know how far AI can go without expensive sensors, custombuilt systems, or massive datasets.

Conclusions on Benefits and Advice for the Sector

The study confirms that turf maintenance is a complex, multivariable task. Playability and turf stress levels are influenced by many interacting factors: mowing frequency, watering schedules, nutrient applications, local weather conditions, and more. For an AI model to provide useful predictions or advice, it must understand and integrate all of these elements.

Unfortunately, the ADORE project found that current state-of-the-art AI time series technology cannot make reliable predictions based on the sparse and irregular data that potentially could be collected manually at most golf courses. Even simple outcomes—like how fast a green will play (measured by the STIMP meter)—could not be predicted with confidence. AI models failed to outperform simple baseline assumptions when trained on a limited amount of manually collected data and weather data from a nearby station. In short, artificial intelligence has the potential to support greenkeeping, but only if it's given enough of the right information.

That said, it shows that AI-based approaches are not entirely out of reach - especially when simplifying and implementing a smarter starting point, which presumably satisfies practical use-cases. For example, AI worked better when predicting binary outcomes, like whether a disease would appear or not, rather than predicting continuous measurements like turf stress levels or play speed. Simplifying the model goals made results more promising, but not yet practically useful.

For clubs considering AI tools in the future, the advice is clear: do not rush into large-scale sensor investments without first identifying what types of data are actually necessary. Instead, the sector should pursue research that investigates which data points have the greatest impact on prediction accuracy. By doing this iteratively, golf clubs can avoid unnecessary spending while still moving toward a future of smarter turf management.

Additionally, the project highlights that expert knowledge from greenkeepers should remain central in the decision-making process. AI should not be seen as a replacement, but as a complement - something that enhances rather than replaces human expertise.

Results

The project utilized four years of uninterrupted weekly data from Golf Club ’t Zelle in the Netherlands. This dataset focused on a specific green and included detailed logs of ongoing maintenance activities, along with occasional soil test results. To complement this, weather data was integrated from the nearest weather station via an open-access weather API.

The research team experimented with multiple types of AI models, including advanced time-series algorithms like LSTM (Long Short-Term Memory), GRU (Gated Recurrent Units), and Temporal Convolutional Networks. These models were trained to simulate various turf condi-tions based on maintenance inputs and environmental data. Key targets included:

- The STIMP value (speed of the green)
- The appearance of diseases like dollar spot and *Fusarium*
- Infestations from larvae (e.g., *Tipulidae*)

At first glance, some results looked promising. For example, certain models achieved up to 100% accuracy in predicting whether a disease would occur. STIMP values, too, were predicted with up to 95% accuracy in some scenarios.

However, a deeper look revealed serious flaws. The models performed well only on very small, specific subsets of data. Their success often came from lucky guesses rather than genuine understanding of patterns. Many models failed to generalize - meaning they could not reliably predict new data they hadn’t seen before.

In fact, when the data was tested on artificially created input scenarios or slightly different greens, the AI models gave inconsistent results. This suggests they had not actually “learned” the true relationships between inputs and outcomes. Furthermore, key factors that should have mattered—like rainfall, temperature, or mowing height—were not always present among the most important features selected by the AI. Even more surprisingly, the leading models did not incorporate their own predictions as input features. For example, models pre-dicting dollar spot disease did not consider whether dollar spot had already been present in previous weeks - a factor that should be highly relevant to the forecast.

For example, in predicting disease outbreaks, models sometimes prioritized irrelevant inputs while ignoring weather data that should logically have a strong impact. Similarly, for STIMP predictions, models required data from up to nine weeks prior- an unusually long lag that doesn’t align with how quickly turf conditions typically change. The implication is that models may have been overfitting - learning quirks in the dataset rather than useful, generalizable rules.

For certain modeling tasks, and depending on the nature of the problem and the data, having only about 20 weeks of testing data per target makes it relatively easy for AI models to achieve high accuracy by chance — especially when thousands of models are evaluated. Some models simply guessed the majority class—predicting “no disease” every time—and still scored high because disease occurrences were rare.

While the results confirm that AI can recognize some patterns, they also raise red flags about the reliability and interpretability of those results under current data limitations.

Methods – Briefly Explained

To test whether AI can help in turf management, the research team trained several AI models using historical data on turf conditions, maintenance routines, and weather. These models are designed to spot patterns over time - essentially learning how today’s decisions and conditions might affect turf tomorrow.

The models used were primarily time-series neural networks, which are well-suited for data that changes over time. The team explored various configurations of input features, model types, and training algorithms, running thousands of experiments to find the most accurate combinations.

The main idea was to predict either changes in turf quality (such as faster or slower greens) or the likelihood of stress conditions like disease or larvae presence. In some cases, the models attempted regression (predicting numerical values like the exact STIMP reading), and in other cases, classification (predicting categories like “disease present” or “disease absent”).
To improve the data, it was enhanced with additional weather information from Koninklijk Nederlands Meteorologisch Instituut (KNMI) API and processed to calculate useful features like “days since last irrigation” or “last nitrogen application.”

Despite this effort, model performance was inconsistent. Many of the models that appeared to succeed were actually just matching training data quirks rather than discovering general truths.

Final Reflections and Future Outlook

Although the ADORE project did not produce a functional AI tool for turf maintenance, it provided a key insight: greenkeepers who commit to four years of weekly data collection should not expect reliable decision support from AI time-series modeling alone.

For golf courses with limited resources, off-the-shelf AI tools for turf simulation and planning are not yet a practical solution. These models require more comprehensive and higher-quality data than most clubs currently collect. However, the potential is evident - AI is already making a tangible impact at courses equipped with dense networks of sensors and weather stations. By refining data requirements and tailoring them to the specific conditions and geography of each course, it becomes possible to define the minimum viable investment in data collection hardware. This, in turn, can lower the barriers to effective AI implementation.

The next steps should focus on defining the minimum data needed for reliable AI predictions. Rather than starting with expensive equipment, research should explore which data points provide the biggest improvement in prediction quality. For instance, is local temperature more important than soil moisture? Is it more useful to track nutrient application timing or mowing frequency? These are the types of practical questions that future studies can answer.

Long-term, AI could help clubs tailor their maintenance to specific greens, respond proactively to turf stress, reduce chemical use, and optimize play conditions. But to get there, the sector needs patience, cooperation between data scientists and turf experts, and a careful, iterative approach to testing and validation.

ADORE serves as a cautionary tale - but also as a roadmap. It reminds us that AI is not a shortcut to perfect turf. It’s a tool, and like any tool, it works best when paired with knowledge, experience, and the right materials.