Using a counterfactual prediction framework to assess site-specific nitrogen fertiliser decisions and trade-offs between production, profit, and emissions at the within-field scale.

Mikaela J Tilse[1*], Sally Poole[1], Patrick Filippi[1], Thomas FA Bishop[1, 2]

[1] Precision Agriculture, Hydrology & Geoinformation Science Laboratory; School of Life & Environmental Sciences; Sydney Institute of Agriculture; The University of Sydney; Eveleigh NSW 2015, Australia.

[2] Sydney Informatics Hub; The University of Sydney; Newtown NSW 2042, Australia.

[*] mikaela.tilse@sydney.edu.au

A paper from the Proceedings of the
17th International Conference on Precision Agriculture and 11th Brazilian Congress on Precision and Digital Agriculture
13-15 July 2026
Porto Alegre, Brazil

Abstract.

Nitrogen (N) fertiliser is a major contributor to greenhouse gas emissions in grain production, but is also essential for crop growth and farm productivity. This study examines the trade-off between N application, production outcomes, and emissions intensity [EI; kg carbon dioxide equivalent (CO_2-e) per tonne (t) of grain yield], and evaluates whether site-specific, variable-rate N management can balance these competing objectives. A commercial winter wheat field in Western Australia was used as a case-study, where an on-farm N response field-scale strip trial in 2023 with three N rates [low (0), field (65), high (130) kg N per hectare (ha^{-1})] revealed spatially variable responses in yield and grain protein content (GPC). A two-stage machine learning approach was used to model the hierarchical relationships between N, yield, and GPC. Counterfactual predictions were then generated to simulate site-specific responses to alternative N rates (0 – 140 kg N ha^{-1}), enabling the evaluation of four N management scenarios: (1) observed as-applied N, (2) simulated grower standard practice (uniform 45 kg N ha^{-1}), (3) N optimized to maximise partial profits, and (4) N optimized to minimise EI. Minimising EI resulted in considerably lower N rates, suggesting reliance on residual soil N, although this may have long-term implications for the crop rotation. Both the profit and EI-optimised strategies better-aligned N inputs with the very dry (rainfall decile 1) seasonal conditions experienced. Notably, the profit-optimised strategy also reduced EI relative to both observed and uniform N applications, demonstrating that economically optimal N management can deliver environmental co-benefits. Spatial mapping of counterfactual predictions for different management scenarios and associated emissions provide further insight into within-field variability. Future work will extend this approach across more fields and seasons to improve decision-support.

Keywords.

Crop yield, grain quality, grain protein content, on-farm experimentation, economically optimum nitrogen rate, greenhouse gas emissions.

Introduction

Nitrogen (N) is a key driver of wheat grain protein content (GPC) and yield, and its management is a key agronomic decision that affects crop yields, grain quality, and farm profitability. Yield gap analysis by Hochman and Horan (2018) suggests that suboptimal N management is a key contributor to the yield gap, which is the difference between actual and attainable yields, in Australian grains, with the average N fertiliser application to dryland cereals and oilseeds being 45 kg N ha^{-1} (Angus and Grace 2017). Synthetic N fertiliser is also a key contributor to on-farm greenhouse gas (GHG) emissions and is one of the largest variable input costs for grain production in Australia (Agriculture Victoria 2017). Increasing climatic and economic instability, as well as growing environmental concerns and market demands for sustainable, low emission or carbon neutral products, are putting pressure on growers to optimize N inputs for improved agronomic, economic, and environmental outcomes. While growers could apply more N to close the yield gap and meet yield potentials, there is a trade-off by increasing the associated total emissions (TE) and potential costs of applying more N fertiliser. Overall, there are tensions between the application of synthetic N fertiliser to achieve grain production and grain quality targets, financial costs, and the reduction of environmental impacts from N fertiliser applications that must be balanced and accounted for.

The TE from grain crops are often reported at the farm or regional scale as kg carbon dioxide equivalent (CO_2-e), and include on-farm (scope 1), energy (scope 2), and off-farm (scope 3) emissions. In Australia, emissions can be calculated using the Greenhouse Accounting Frameworks (GAF), which are Australian-specific GHG accounting frameworks developed for a range of industries, including cropping, by the Primary Industries Climate Challenges Centre (PICCC), and align with Australia's National GHG Inventory (Lopez *et al*. 2024). While TE provides a measure of all emissions associated with production, emissions are more commonly compared in terms of emissions intensity (EI), expressed as kg CO_2-e per tonne (t) of grain yield, as this accounts for differences in productivity. This distinction is important as most industry and policy targets (e.g. National Farmers Federation 2030 Roadmap, NFF 2019; Australia's Net Zero Plan, DCCEEW 2025) are framed around reducing EI rather than absolute TE.

Total emissions can vary considerably season-to-season depending on weather and N inputs, due to the risk of volatilization and other N losses, as well as differences in crop production levels, particularly emissions associated with stubble. For example, seasons that experience droughts often have the lowest TE due to reduced input use, leading to less N loss and lower biomass production (Sevenster *et al*. 2022). However, this does not necessarily translate into improved EI, as lower yields may offset reductions in emissions. Given that crop yields can vary considerably within and between fields, farms, and seasons due to differences in soil type, management, and seasonal conditions, and inputs are often variably applied to respond to these differences, both TE and EI are expected to vary considerably across a range of spatial and temporal scales. However, emissions are typically calculated at the farm/enterprise level using the GAF, meaning that within-field variability in TE and EI is currently not considered.

Precision agriculture, in particular site-specific crop management (SSCM), provides a pathway for growers to tailor inputs to the localized requirements of the crop or soil. The SSCM of N fertilization is challenging for growers because of difficulties in accurately mapping soil, and more specifically, inorganic N content, and the significant within-field and seasonal variability observed in Australian cropping systems (Ray *et* al. 2015). Traditionally, N fertiliser recommendations are based on a combination of soil tests, N mass-balance equations, and regional N-response field trial experiments (e.g. Meisinger *et al*. 2008, Meier *et al*. 2021). However, these do not account for localized, within-field variability in N requirements or crop responses to soil, seasonal, or management conditions.

Field trials and on-farm experimentation (OFE) can provide valuable, context-specific insights into crop responses under field conditions by evaluating outcomes from different inputs or management decisions (Tanaka *et al*. 2026). By applying adjacent strips of different fertiliser rates (e.g. low/zero, grower-standard, high/rich), on-farm N-response strip trials can be used to understand localized responses of the crop across different soil types and seasonal conditions to estimate optimum N rates (e.g. the economically optimum N rate; EONR). However, analysis methods such as moving window quadratic response curves (e.g. Colaço et al. 2024) are typically restricted to the trial area, making it difficult to extrapolate N decisions spatially across a field and limits the ability to generate site-specific recommendations.

To address this limitation, we use the relationship observed from an on-farm N-response trial to make predictions about how a winter wheat crop would have responded to different N rates at discrete locations across a field. We refer to these predictions as "counterfactual" because they represent outcomes under N rates that differ from those that were actually applied, but the underlying approach is simple. It involves:

1) Fitting a predictive model using total applied N and observed yield and GPC responses, and
2) Substituting new candidate N values into this fitted model to estimate hypothetical yield and GPC outcomes for alternate N rates.

This produces counterfactual predictions which explore "what-if" scenarios, or what could have happened if a different N rate had been applied. These predictions are retrospective and explore hypothetical outcomes for the season in which the trial was conducted.

Using data derived from an on-farm N-response trial in one case-study field, we demonstrate a practical application of this counterfactual approach to provide spatially explicit predictions of yield and GPC under different N application scenarios. Further, we calculate the TE and EI associated with each N scenario for each location within the field. Results from four N application scenarios were compared:

1) Observed as-applied N application, including the strip trial;
2) Simulated grower standard practice (GSP) of uniformly applying 45 kg N ha^{-1} (i.e. the national average, Angus and Grace 2017);
3) Simulated optimum N rate that maximises partial profit (i.e. the EONR); or
4) Simulated optimum N rate that minimises EI.

For all four scenarios, yield, GPC, and EI were predicted across the case-study field and results were mapped spatially.

Methods

On-farm nitrogen-response experimental trial

To generate yield and GPC responses to different applied N fertiliser rates, an experimental N-response trial was conducted in a 100 ha case-study field for the 2023 growing season in Western Australia, approximately 250 km north of Perth. The farm has a Mediterranean climate (Beck *et al*. 2018) with cool, wet winters and hot, dry summers and an annual average rainfall of 493 mm (BOM 2024). For 2023, rainfall was very low, totaling 248 mm for the year (decile 1). The dominant soil type across the farm is Tenosols according to the Australian Soil Classification Map (Searle 2021).

The trial consisted of adjacent strips receiving high, field, and low (zero) N rate treatments applied over two applications, based on the trial design from Colaço *et al* (2024). Urea Ammonium Nitrate (UAN; 32% N) was applied at sowing at rates of 90 (high), 45 (field) and 0 (zero) liters (L) UAN/ha, followed by a mid-season top-up of urea (42% N) at rates of 200 (high), 110 (field), and 0 (zero) kg urea/ha. Each strip was 36 m wide, equivalent to three combine-harvester widths. Management zones within the field were delineated using k-means clustering with available spatial data layers including historical yield maps, gamma radiometric data, electromagnetic (EM) induction data, and elevation derived from the Geoscience Australia Shuttle Radar Topography Mission (SRTM)-derived 1 arc second digital elevation model (DEM; Gallant et al. 2009). These management zones were used to determine the trial placement within the field in order to capture sufficient variability in crop response. The remainder of the field was managed according to standard grower practice, with uniform UAN application (45 L UAN/ha) at sowing, and variable urea application between zones mid-season (90, 100, and 120 kg Urea/ha).

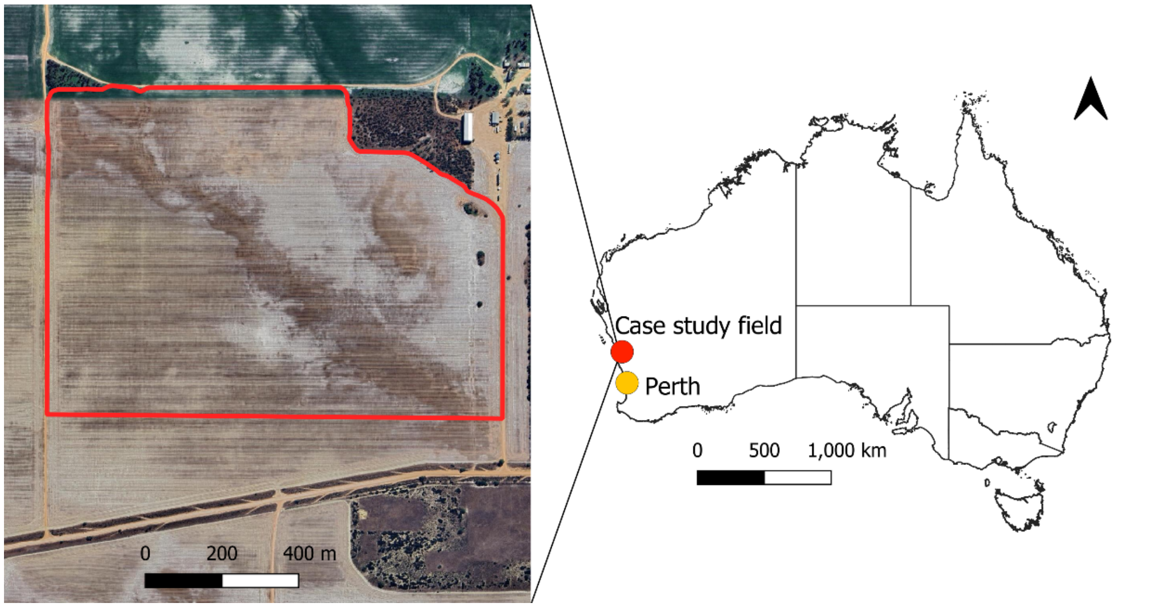

Figure 1. Location of the case-study field in Western Australia.

Wheat grain yield and protein data were collected for the case-study field using John Deere S790 combine harvesters equipped with the ActiveYield^{TM} automatic yield calibration system, and the HarvestLab 3000^{TM} near-infrared (NIR) spectroscopy sensor for GPC (Tilse et al. 2024). Approximately one measurement was performed per second (Schade et al. 2023), with a harvester header width of 12 m. All yield and protein data was cleaned as per the methods from Taylor et al. (2007) and to values recommended by Taylor et al. (2007), Blakeney et al. (2009) and Žilić et al. (2011). Yield values < 0 and > 10 tonnes per hectare (t/ha), and protein values < 8 and > 22%, were removed. Measurements more than 2.5 standard deviations (SDs) above and below the field mean were also removed, and concurrent duplicate protein values were excluded based on their time and location to account for data errors during harvest. The yield and protein data was then rasterised onto a 10 m grid. To preserve the integrity of the N-response strip trial, interpolation was performed separately for each treatment strip and the surrounding non-trial area. A local k-nearest neighbour approach was used in each strip due to its narrow width (36 m), while inverse distance weighted kriging was applied to the remaining field area to represent broader spatial trends. The resulting surfaces were then combined into a single raster for subsequent analysis. All data processing and analysis was performed using R Statistical Software (v4.3.1; R Core Team 2023).

Greenhouse gas emission calculations

The TE and EI were calculated by adapting the Grains GAF (G-GAF) tool developed by Lopez et al. (2024) [version 11.0] into R. The TE were calculated as the sum of the scope 1 (on-farm), 2 (energy), and 3 (off-farm) emissions, which were calculated on a per-hectare basis in kg CO_2-e ha^{-1} for each pixel within the field. The EI represents the TE (kg CO_2-e ha^{-1}) per unit (tonne) of grain produced (presented as kg CO_2-e t^{-1}). Both the TE and ET were calculated for each pixel, meaning that they could also be mapped within the case-study field. This study assumed that the applied fertiliser, yield, and associated crop residue were the only dynamic variables affecting emissions, and that all other emissions were consistent within the field. The following assumptions and generalisations were made about the crop to determine the TE:

- None of the fields' crop residue was burnt for stubble management.
- Trees sequestered no carbon because the study only focuses on the within-field GHG emissions of the crop.
- No petroleum, LPG, or electricity was used in the production of the crop, and the average diesel consumption was 11 L ha^{-1} (Gjerek *et al*. 2021).
- No soil amelioration with lime had occurred.
- The amount of general pesticides and specific herbicides (Paraquat, Diquat, Glyphosate) was calculated based on a typical wheat spray program.

Machine learning for grain yield and protein content modelling

A series of eXtreme Gradient Boosting (XGBoost) models were built to predict yield and GPC separately using the 'xgboost' R package (v1.7.7.1; Chen et al. 2024). XGBoost models are a decision-tree-based machine learning algorithm that uses gradient boosting to build an ensemble of trees for prediction. Boosting is applied to weak learners (regression trees) sequentially so that each subsequent learner improves the errors of the previous learners, resulting in a stronger model that minimises training errors (Nielsen 2016; Fauzan and Murfi, 2018).

Machine learning models like XGBoost can handle complex datasets with many predictor variables, enabling us to capture the impacts of these different factors (Liakos et al. 2018; Rudin 2019), but simple models using fewer covariates are easier to interpret. For this study, three features were used to model wheat grain yield [total applied N, the soils apparent electrical conductivity (EC_a), and elevation derived from a DEM], and four features were used to model wheat GPC (total applied N, the soils EC_a, elevation derived from a DEM, and observed yield). These features were chosen as they are easy to interpret and make agronomic sense as they can (in)directly describe the factors influencing variability in yield and GPC, based on known relationships from previous work conducted in this case-study field (Poole *et al*. 2024; Tilse *et al*. 2024). Briefly, elevation derived from the Geoscience Australia Shuttle Radar Topography Mission (SRTM) derived 1 arc second DEM (Gallant et al. 2009; accessed via the 'dataharvester' R package, v0.1.2; Harianto et al. 2023) was used to describe elevation change within each field and represents the patterns of water flow and accumulation. The soils EC_a, which was measured and mapped across each field with an electromagnetic (EM) induction survey, was used as a proxy to describe relative variability in soil texture and water holding capacity. Yield was used to represent the cumulative effects of all other seasonal conditions experienced by the crop, such as management decisions, weather, and inherent soil properties (i.e. naturally occurring characteristics that are not easily altered, such as soil depth). Total applied N was calculated by determining the percent composition of N per unit of product (UAN = 32% N, Urea = 46% N) and calculating the total applied N across the season in kg N/ha. The TE and EI, elevation, total applied N, and soil EC_a values were then extracted onto the same 10 m grid used for the yield and protein data to form a single dataframe/datacube.

A two-stage modelling approach was used to represent the hierarchical relationship between N, grain yield, and GPC. Nitrogen influences both yield and GPC. The GPC is essentially a measure of N concentration in the grain, and is partially determined by yield through a dilution or concentration effect (Terman *et al*. 1969; Simmonds 1995). Therefore, the relationship between N and GPC is mediated, in part, by yield and these components are not independent.

First, an XGBoost model was used to build a predictive model of grain yield using total applied N, soil EC_a, and elevation as input features (Equation 1).

$$\text{Yield} \sim \text{elevation} + EC_a + \text{total applied N} \tag{1}$$

A second XGBoost model was then used to build a predictive model of GPC using observed yield, total applied N, soil EC_a, and elevation as input features (Equation 2). Including yield as a predictor allows the model to capture the dilution effect, where, as the amount of carbohydrates in the grain (yield) increases, GPC decreases (Terman *et al*. 1969; Simmonds 1995), as well as the indirect influence of N on GPC through its effect on yield (Holford *et al*. 1992; Pan *et al*. 2020).

$$\text{Grain protein content} \sim \text{Yield} + \text{elevation} + EC_a + \text{total applied N} \tag{2}$$

This two-stage approach enables a separation of the influence of N on yield, N on GPC, and the intermediate effect of yield on GPC, as driven by N.

Both models were evaluated using a 10-fold cross-validation, and model quality was evaluated based on the Lins concordance correlation coefficient (LCCC; Lin 1989), which is a measure of the agreement between observed and predicted values in relation to the 1:1 line where 1 is a perfect fit, and the root mean square error (RMSE), which is a measure of the difference between observed and predicted values where the smaller the RMSE, the better.

Using counterfactual predictions to make site-specific nitrogen recommendations

Modelling the yield and GPC responses to a range of N rates induced through the OFE trial (0 – 130 kg N ha^{-1}; Equations 1 and 2) enabled the fitted models to capture the observed relationships between N, yield, and GPC under the seasonal conditions and inherent soil/landscape features of the case-study field. These fitted relationships were then used to predict changes in yield and GPC under alternative N application rates.

For a range of candidate N rates (0 – 140 kg N/ha at 5 kg intervals; n = 28), an updated yield prediction was made across the field by substituting the total N (Equation 1) with each candidate N into the XGBoost model (Equation 3). These updated yield predictions capture how yield responds to different candidate N rates based on the modelled relationships in the observed data.

$$\text{Updated Yield} \sim \text{elevation} + EC_a + \text{total applied N} \tag{3}$$

Then, for each of these candidate N rates, an updated GPC prediction was made by substituting the total N (Equation 2) with each candidate N and the corresponding updated yield prediction (from Equation 3) into the XGBoost model (Equation 4). This new GPC prediction accounts for the interconnections between yield, GPC, and N by capturing how GPC responds to different N rates, while also accounting for how yield responds to N and the impact of yield on GPC.

$$\text{Updated Grain Protein Content} \sim \text{updated Yield} + \text{elevation} + EC_a + \text{total applied N} \tag{4}$$

The associated partial profit, TE, and EI were then calculated at each pixel within the field for every candidate N rate, based on the updated yield and GPC predictions from equations 3 and 4. The partial profit was calculated using gross income minus expenditure on fertiliser inputs. While both UAN and urea were applied to the case-study field, urea was used to estimate fertiliser input costs to calculate the EONR that maximised partial profit. Gross income was paid based on the GPC grade [GPC < 10.5% = $261.8/t, GPC 10.5 – 11.5% = $308/t, GPC 11.5 – 13% = $323.4/t, > 13% = $338.8/t; Colaço et al. 2024] and urea was priced at $700/t.

Based on the counterfactual predictions, the applied N rate, yield, GPC, and EI were calculated within each field for four N application scenarios:

1) Observed as-applied N application, including the strip trial;
2) Simulated GSP of uniformly applying 45 kg N ha^{-1};
3) Simulated optimum N rate that maximises partial profit (i.e. the EONR); or
4) Simulated optimum N rate that minimises EI.

Results and discussion

Yield, protein, and emissions responses to on-farm nitrogen-response trial

There was substantial variability observed in grain yield, ranging from 0 – 2.6 tonnes ha^{-1} (Figure 2b), and GPC, ranging from 11.4 – 17.5% (Figure 2c). There was a strong inverse relationship observed between yield and GPC at the whole-field-scale (Pearsons correlation = -0.64), as expected (Terman *et al*. 1969; Simmonds 1995). Grain protein responded much more strongly to the different applied N rates in the trial strips compared to yield.

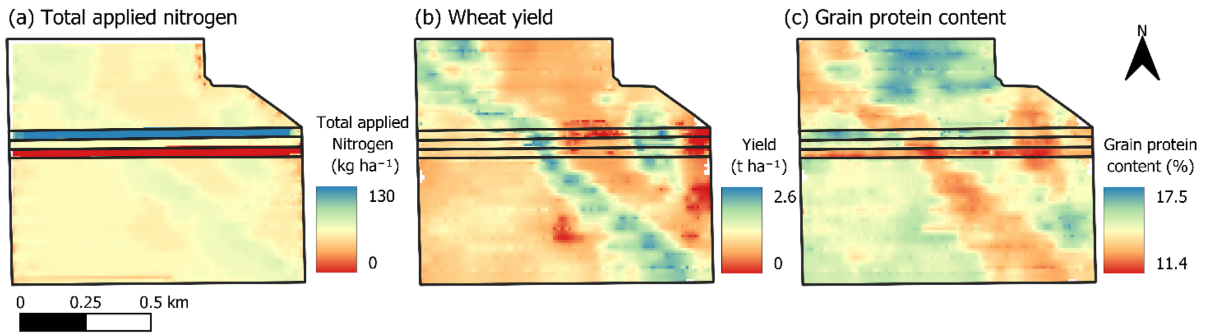

Figure 2. Total applied nitrogen (kg per hectare, ha-1; a), observed wheat grain yield (tonnes per hectare, t ha-1; b) and observed grain protein content (%; c) across the case-study field for the 2023 season.

A tension exists between the application of additional inputs, including N fertilisers, to increase crop production, and the contribution of these additional inputs to increased emissions. The TE were lowest under the zero applied N strip, with emissions in the strip arising from crop residues, pesticides, and fuel use from field operations (Figure 3). The TE substantially increased with increasing rates of applied N, as observed in the high applied N strip. While higher rates of applied N fertiliser increase TE, if these additional inputs contribute to increased crop production, meaning that the ET is lower in those parts of the field. In this case study field, the map of EI indicates that EI is lower where yield is higher – particularly in the distinct diagonal soil-landscape feature running north-west to south-east through the paddock. The EI is also low where applied N is low but the crop still yielded, particularly throughout the zero-N strip. This highlights the trade-off between the emissions and productivity.

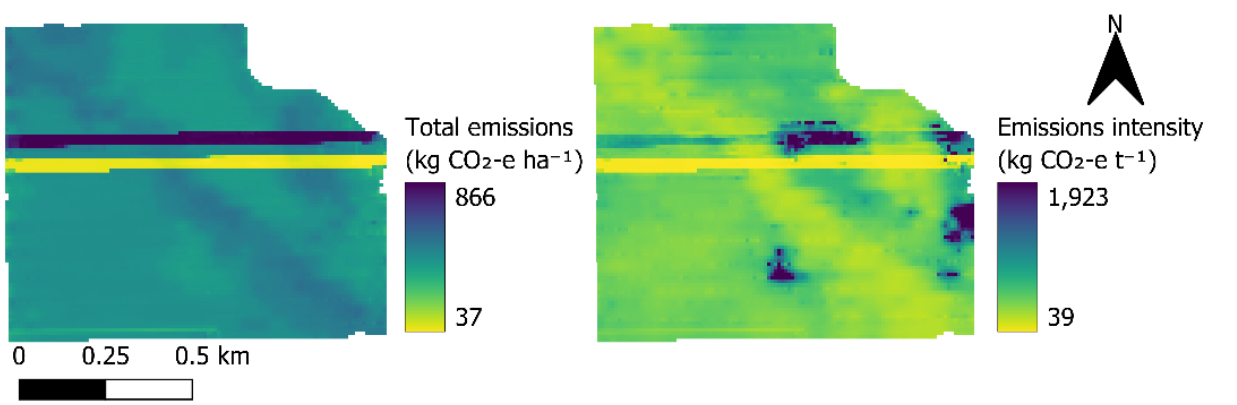

Figure 3. Observed total emissions per hectare (ha) (kg carbon dioxide equivalent per hectare, CO_2-e ha^{-1}) and emissions intensity per tonne (t) of grain per ha (kg CO2-e t^{-1}) across the case-study field for the 2023 season.

Most of the work on calculating emissions is done at the farm/enterprise level (e.g. d'Abbadie and Kingwell 2025; Khaqan *et al*. 2026). In this work, we adapted the G-GAF tool (Lopez et al. 2024) to calculate emissions at the pixel level, enabling us to map emissions spatially at the within-field scale (Figure 3). This approach utilises existing on-farm datasets, including maps of grain yields, GPC, and as-applied nutrients, meaning that this approach can be adopted across any field with these available data layers. While farm/enterprise level assessments can identify general trends and sources of emissions, and potentially identify areas for improvement or optimisation, understanding variation at the within-field level can highlight particular regions of a field/farm where management can be optimised. For example, areas with a high EI (Figure 3) may warrant further investigation to identify the drivers of poor yield responses, or a calculation of the optimum N rate to better match inputs to inherent soil features.

Model quality evaluation

For the case-study field, results from the 10-fold cross-validation of the XGBoost model for yield prediction produced an RMSE of 0.15 t/ha, and an LCCC of 0.95. For GPC, the XGBoost model produced an RMSE of 0.29% and an LCCC of 0.96. This shows that both yield and GPC are well represented by the covariates within the case study field. This is not surprising given that in Western Australia, spatial variability in wheat grain yield and protein content within-fields is predominately driven by variation in plant-available water capacity (Lawes *et al*. 2009), and previous research conducted across this case study farm has indicated that N and soil moisture are the key drivers of variability (Poole *et* al. 2024; Tilse *et al.* 2024). The application of this counterfactual approach over more regions and seasons, and the assessment of different covariates, should be explored in future work. Depending on the regional context, more complex models may be necessary to capture interactions between other genetic, environmental, and management factors.

Using counterfactual predictions to make site-specific nitrogen recommendations

The counterfactual prediction framework was used to estimate yield, GPC, and EI outcomes under four different N application scenarios. For the applied N rate (Figure 4 a – d) that achieved each of the four scenarios [1) Observed, 2) GSP, 3) Maximise partial profit, 4) Minimise EI], the resulting yield (Figure 4 e – h), GPC (Figure 4 i – l), and EI (Figure 4, m – p) from the counterfactual predictions were mapped. Importantly, these scenarios are retrospective and quantify what would have occurred under alternative N application strategies, given the observed relationships between N, yield, and GPC in the 2023 season. As such, the predictions reflect hypothetical outcomes conditional on the specific seasonal and environmental conditions experienced in that year. The summary statistics for the four scenarios are also presented in Table 1.

Table 1. Mean and standard deviation (SD) of applied nitrogen [N, kg N per hectare (ha)$^{-1}$], yield [tonnes (t) ha^{-1}], grain protein content (GPC, %), and emissions intensity (EI, kg carbon dioxide equivalent per t of grain, kg CO_2-e t^{-1}) under four N application scenarios: 1) observed as-applied N application (including the strip trial), 2) simulated grower standard practice (GSP) of uniformly applying 45 kg N ha^{-1}, 3) simulated optimum N rate that maximises partial profit, and 4) simulated optimum N rate that minimises EI.

	Applied N (kg N ha^{-1})		Yield (t ha^{-1})		GPC (%)		EI (kg CO_2-e t^{-1})	
Scenario	**Mean**	**SD (±)**	**Mean**	**SD (±)**	**Mean**	**SD (±)**	**Mean**	**SD (±)**
1) Observed	64.37	19.78	1.13	0.47	14.53	1.08	441.32	198.26
2) GSP	45.00	0.00	1.12	0.28	14.60	0.77	325.72	93.84
3) Maximum profit	33.51	29.93	1.50	0.41	13.15	0.63	176.42	115.00
4) Minimum EI	0.14	2.62	1.16	0.36	12.83	0.58	60.28	19.98

Overall, mean yield was constant across all scenarios except for scenario 3 (maximise partial profit), in which yields increased. Under scenario 4 (minimise the EI), the applied N recommendation was considerably lower compared to all other scenarios. While this suggests that the EI could be minimised by relying on soil N, mining the soil N reserves would likely have ongoing consequences for crops in the following seasons (Hochman and Horan 2018). Furthermore, this particular growing season had very low rainfall (decile 1) and it is more likely that the crop was moisture-limited, rather than N-limited. The applied N recommendations for scenario 3 (maximise partial profit) and scenario 4 (minimise EI) were both better-aligned with the actual seasonal conditions experienced and resulted in the avoidance of unnecessary fertiliser inputs. In seasons with average or above-average rainfall, stronger yield responses to applied N would be expected, and N recommendations based on maximising profits Likewise, given that the EI is a trade-off between total emissions and the amount of grain produced, it is possible that in a higher rainfall year, EI could still have been minimised with higher applied N rates if there was a stronger yield response to applied N. It is also interesting to note that the EI under scenario 3 (maximise partial profit) was considerably lower than under scenario 1 (observed as-applied N application) and scenario 2 (simulated GSP of uniformly applying 45 kg N ha^{-1}; Table 1). This highlights that even without prioritising emissions reductions, reallocating N inputs to maximise partial profits (which accounts for both yield and the grain quality i.e. GPC) may still contribute to improved environmental outcomes as well.

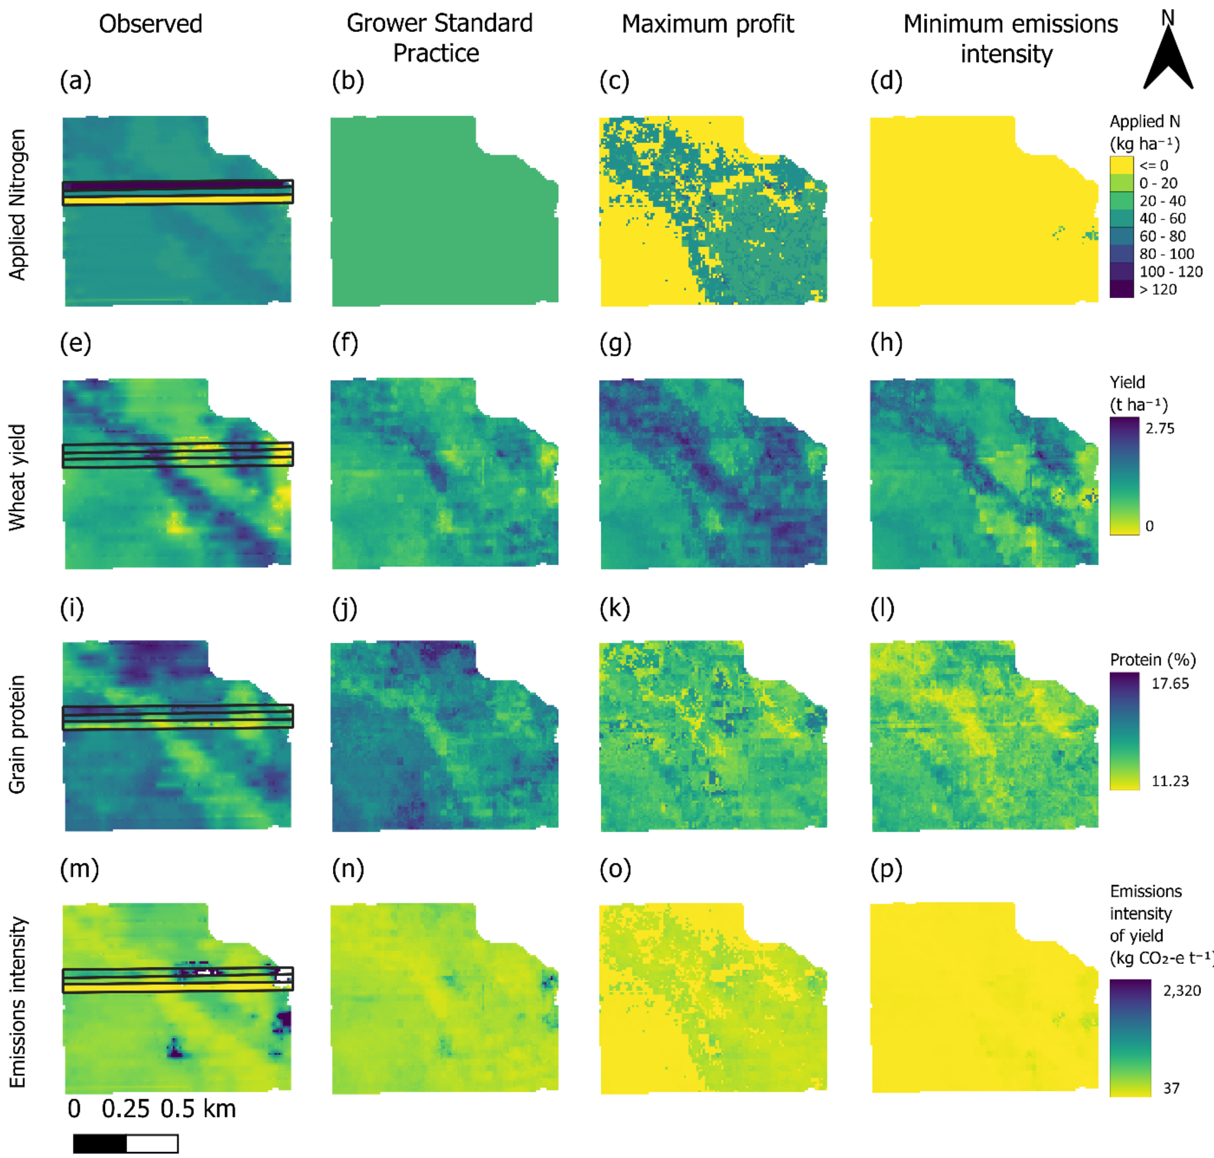

Figure 4. Applied nitrogen (kg N per hectare, kg N ha^{-1}), grain yield [tonnes (t) ha^{-1}], grain protein content (%), and emissions intensity (EI, kg carbon dioxide equivalent per t of grain, kg CO_2-e t^{-1}) under four N management scenarios: 1) observed as-applied N application (including the strip trial), 2) simulated grower standard practice (GSP) of uniformly applying 45 kg N ha^{-1}, 3) simulated optimum N rate that maximises partial profit (i.e. the EONR), and 4) simulated optimum N rate that minimses EI.

The applied N recommendations provided here are retrospective based on experimental data already observed. The retrospective recommendations under scenarios 3 and 4 represent considerable potential gains compared to what was applied by the grower (scenario 1) or GSP based on national average fertiliser use (scenario 2). Yet in practice, fertiliser decisions are made prior-to or during the season based on best-available information, including uncertain rainfall forecasts and soil moisture availability estimates. While responses will vary season-to-season, if more experimental trials are conducted over multiple seasons, we can begin to build a database to explore potential N strategies depending on the seasonal outlook. For example, if it is expected that the upcoming season will be wetter than average, results from previous experimental trials conducted in similar seasons within the same field could be used to explore potential impacts of different N strategies given the seasonal conditions at hand. Future work should seek to conduct and analyse more field-trials and OFE across a range of regions and seasons to improve N fertiliser decision-making.

While the current emissions reporting mechanisms focus on-farm/enterprise level assessments, the counterfactual prediction framework explored in this study provides growers with a greater understanding of the spatial variability of emissions within fields, and an opportunity to explore how different N management strategies could affect both production outcomes and associated emissions. This framework could support growers to develop management strategies that achieve emissions reductions within fields and across farms by implementing more sustainable SSCM strategies and increasing climate-smart, sustainable intensifications of their farming systems. For example, the reallocation of N inputs to areas with a stronger yield response may be advantageous to increasing yields and return on investment, as well as minimising the EI. Overall, understanding the link between management and emissions at the within-field and farm scales can enable growers to improve economic and environmental sustainability outcomes, ultimately helping to future-proof their enterprises.

Conclusion

This study demonstrates how counterfactual prediction frameworks, informed by precision agriculture data and OFE, can be used to evaluate site-specific N management decisions and associated trade-offs between yield, GPC, profitability, and EI. By extrapolating observed relationships between N, yield, and GPC from a strip trial to the whole field using a two-stage machine learning modelling framework, this approach enables the retrospective evaluation of alternative N strategies beyond direct treatment comparisons. Results from this case study showed that both profit- (scenario 3) and EI-optimised (scenario 4) strategies improved outcomes relative to observed (scenario 1) and uniform GSP (scenario 2) practices. Notably, optimising for profit still resulted in substantial reductions in N and associated EI, highlighting the potential for economic and environmental co-benefits through a better allocation of inputs. The applied N recommendations presented here are retrospective, based on observed seasonal conditions. Expanding this approach across multiple seasons and fields would allow the development of a larger database to support decision-making under different seasonal outlooks. While current emissions reporting frameworks focus on farm/enterprise-scale assessments, this approach provides spatially explicit insights into within-field variability in both production and emissions outcomes. This enables growers to explore how reallocating N inputs to areas with stronger yield responses may improve return on investment while also reducing EI.

Acknowledgments

This research was funded by the Grains Research and Development Corporation (GRDC) project “PrecisionProtein: Mapping Grain Quality for Improved Management” (UOS2507 – 002RTX), and the Christian Rowe Thornett Stipend Scholarship and the Thomas Lawrance Pawlett Scholarship (Supplementary) within the Sydney Institute of Agriculture at The University of Sydney. Soil surveys conducted as part of the GRDC project “3D Plant Available Water Capacity (PAWC) and constraint mapping” (UOS2206-009RTX) provided valuable background information on soil and farm conditions which supported this research. We extend our thanks to the growers, consultants, Precision Cropping Technologies, and John Deere for collaborating, facilitating and providing access to this data.

References

Agriculture Victoria (2017) Summary report: Reducing on-farm nitrous oxide emissions through improved nitrogen use efficiency in grains. The State of Victoria Department of Economic Development, Jobs, Transport and Resources.

Angus JF, Grace PR (2017) Nitrogen balance in Australia and nitrogen use effeiciency on Australian farms. *Soil Research* **55**, pp. 435–450. doi: 10.1071/SR16325.

Beck HE, Zimmermann NE, McVicar TR, Vergopolan N, Berg A, Wood EF (2018) Present and future köppen-geiger climate classification maps at 1-km resolution. *Scientific Data* **5**, pp. 1–12. doi: 10.1038/sdata.2018.214.

Blakeney AB, Cracknell RL, Crosbie GB, S.P. J, Miskelly DM, O. BL, Panozzo JF, D.A.I. S, Solah V, Watts T, Westcot T, Williams RM (2009) *Understanding Australian Wheat Quality*. Grains Research and Development Corporation (GRDC).

Bureau of Meteorology (BOM) (2024) Summary statistics Eneabba. *Climate Statistics for Australian Locations*. Retrieved 15th December 2024 from http://www.bom.gov.au/climate/averages/tables/cw_008225.shtml.

Chen T, He T, Benesty M, Khotilovich V, Tang Y, Cho H, Chen K, Mitchell R, Cano I, Zhou T, Li M, Xie J, Lin M, Geng Y, Li Y, Yuan J (2024) xgboost: Extreme Gradient Boosting. R package version 1.7.7.1. doi: 10.32614/CRAN.package.xgboost.

Colaço AF, Whelan BM, Bramley RGV, Richetti J, Fajardo M, McCarthy AC, Perry EM, Bender A, Leo S, Fitzgerald GJ and Lawes RA (2024) Digital strategies for nitrogen management in grain production systems: lessons from multi-method assessment using on-farm experimentation, *Precision Agriculture* **25**(2), pp. 983–1013. doi: 10.1007/s11119-023-10102-z.

d'Abbadie C, Kingwell R (2025) Western Australian and South Australian Grain Enterprise Emission Intensities and Gross Margins: A Review. *Australian Farm Business Management Journal*, **22**(3). ISSN: 1449–7875.

Department of Climate Change, Energy, the Environment and Water (DCCEEW) (2025) *Australia's Net Zero Plan*.

Fauzan MA., Murfi, H (2018). The accuracy of XGBoost for insurance claim prediction. *International Journal Advance in Soft Computing and Its Application*, **10**(2), pp. 159–171.

Gallant J, Wilson N, Tickle PK, Dowling D, Read A (2009). 3 second SRTM Derived Digital Elevation Model (DEM) version 1.0. Geoscience Australia, Canberra. Available at: https://pid.geoscience.gov.au/dataset/ga/72759.

Gjerek M, Morgan A, Gore-Brown N, Womersley G (2021) Diesel Use in NSW Agriculture and Opportunities to Support Net Zero Emissions. Sydney: Australian Alliance for Energy Productivity for NSW Department of Primary Industries.

Harianto J, Haan S, Butterworth N (2023) dataharvester: Download and Process Geospatial Data. R package version 0.1.2. Retrieved January 8th 2025 from https://sydney-informatics-hub.github.io/dataharvester/.

Hochman Z, Horan H (2018) Causes of wheat yield gaps and opportunities to advance the water-limited yield fronteir in Australia. *Field Crops Research* **288**, pp. 20–30. doi: 10.1016/j.fcr.2018.08.023.

Holford ICR, Doyle AD, and Leckie CC (1992) Nitrogen response characteristics of wheat protein in relation to yield responses and their interactions with phosphorus. *Australian Journal of Agricultural Research*, **43**(5), pp. 969–986. doi: 10.1071/AR9920969.

Khaqan K, Fest B, Harrison T, Sandhu H (2026) Carbon Accounting at the Farm Level: Scenarios to Achieving Net-Zero Status. *Journal of Sustainable Agriculture and Environment*, e70180. doi: 10.1002/sae2.70180.

Lawes RA, Oliver YM, Robertson MJ (2009) Integrating the effects of climate and plant available soil water holding capacity on wheat yield. *Field Crops Research* **113**. doi: 10.1016/j.fcr.2009.06.008

Liakos KG, Busato P, Moshou D, Pearson S, Bochtis D (2018) Machine learning in agriculture: A review, *Sensors* **18**(8), 2674. doi: 10.3390/s18082674.

Lin L I-K. (1989) A concordance correlation coefficient to evaluate reproducibility, *Biometrics*, **45**(1), 255–268. doi: 10.2307/2532051.

Lopez MB, Ekonomou A, Eckard RJ (2024) A Greenhouse Accounting Framework for crop production (G-GAF) based on the Australian National Greenhouse Gas Inventory methodology. Updated June 2024. 11.0 ed.

Meier EA, Hunt JR, Hochman Z (2021) Evaluation of nitrogen bank, a soil nitrogen management strategy for sustainably closing wheat yield gaps. *Field Crops Research* **261**, 108017. doi: 10.1016/j.fcr.2020.108017.

Meisinger JJ, Schepers JS, Raun WR (2008) Crop nitrgen requirement and fertilization. In: Schepers JS and Raun WR (eds), *Nitrogen in Agricultural Systems*, Agronomy Monographs, Madison. pp. 563–612. doi: 10.2134/agronmonogr49.c14.

National Farmers Federation (NFF) (2019) *2030 Roadmap*. Last accessed on 15 June 2026 at: https://nff.org.au/policies/roadmap/

Nielsen D (2016). *Tree boosting with xgboost-why does xgboost win" every" machine learning competition?* Masters thesis. NTNU (Norwegian University of Science and Technology).

Pan, WL, Kidwell KK, McCracked VA, Bolton RP, Allen M (2020), Economically Optimal Wheat Yield, Protein and Nitrogen Use Component Responses to Varying N Supply and Genotype, *Frontiers in Plant Science*, **10**, pp. 1–10. doi: 10.3389/fpls.2019.01790.

Poole S, Bishop TFA, Al-Shammari D and Filippi P (2024) Mapping within-field drivers of crop production in space and time using interpretive machine learning. In: Lawes R, Flower K, Mwenda G, and Singh B (eds) *Adaptive agronomy for a resilient future. Proceedings of the 21st Australian Society of Agronomy Conference*.

R Core Team (2023) R: A Language and Environment for Statistical Computing, *R Foundation for Statistical Computing, Vienna, Austria.* Retrived January 8th 2025 from https://www.r-project.org/.

Ray DK, Gerber JS, MacDonald GK, West PC (2015) Climate variation explains a third of global crop yield variability. *Nature Communications* **6** 5989. doi: 10.1038/ncomms6989.

Rudin C (2019) Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead, *Nature Machine Intelligence*, **1**(5), pp. 206–215. doi: 10.1038/s42256-019-0048-x.

Schade P, Struve C, Gralfs R (2023). ‘Sensor Technology for Continuous Measurement of Constituents During Small Grain Harvest’. In: *ATZheavy duty worldwide*. **16**(2), pp. 10–15.

Searle R (2021) Australian Soil Classification Map. *Terrestrial Ecosystem Research Network*. Last accessed on 10 October 2025 at: https://portal.tern.org.au/metadata/TERN/15728dba-b49c-4da5-9073-13d8abe67d7c.

Sevenster M, Bell L, Anderson B, Jamali H, Horan H, Simmonds A, Cowie A, Hochman Z

(2022) Australian Grains Baseline and Mitigation Assessment. Main Report. CSIRO, Australia.

Simmonds NW (1995) The relation between yield and protein in cereal grain. *Journal of the Science of Food and Agriculture*, **67**, pp. 309–315. doi:10.1002/jsfa.2740670306.

Tanaka TST, Colaço AF, Mieno T, Riley SS (2026) Advanced data anlytics for on-farm experimentation: methods, challenges, and opportunities for modern agronomic decision making. *Plant Production Science* **29**(2), pp. 271–292. doi: 10.1080/1343943X.2026.2664854.

Taylor JA, McBratney AB, Whelan BM (2007) Establishing management classes for broadacre agricultural production. *Agronomy Journal* **99**, pp. 1366–1376. doi:10.2134/agronj2007.0070.

Terman GL, Ramig RE, Dreier AF and Olson RA (1969) Yield-Protein Relationships in Wheat Grain, as Affected by Nitrogen and Water 1, Agronomy Journal **61**(5), pp. 755–759. doi: 10.2134/agronj1969.00021962006100050031x.

Tilse MJ, Bishop TFA, Filippi P (2024) Understanding the drivers of within-field grain protein content variability with spatial data layers. In In: Lawes R, Flower K, Mwenda G, and Singh B (eds) *Adaptive agronomy for a resilient future. Proceedings of the 21st Australian Society of Agronomy Conference*.

Žilić S, Barać M, Pešić M, Dodig D, Ignjatović-Micić D (2011) Characterization of proteins from grain of different bread and durum wheat genotypes. *International Journal of Molecular Sciences* **12**, pp. 5878–5894. doi:10.3390/ijms12095878.