Forecasting for Sustainable Dairy Produce: Enhanced Long-Term, Milk-Supply Forecasting Using *k*-NN for Data Augmentation, with Prefactual Explanations for XAI

Eoin Delaney[1,2,3], Derek Greene[1,2,3], Laurence Shalloo[3,4], Michael Lynch[5], and Mark T. Keane[1,2,3]

[1] School of Computer Science, University College Dublin, Ireland
[2] Insight Centre for Data Analytics, University College Dublin, Ireland
[3] VistaMilk SFI Research Centre, Ireland
[4] Teagasc, Animal and Grassland Research Centre, Cork, Ireland
[5] ICBF, Cork, Ireland
{eoin.delaney,derek.greene,mark.keane}@insight-centre.org, laurence.shalloo@teagasc.ie, mlynch@icbf.com

Abstract. Accurate milk supply forecasting for the dairy sector, covering 1000s of farms with low resolution data, is a key challenge in achieving a sustainable, precision agriculture that can improve farm management, balancing costs, energy use and environmental protection. We show that case-based reasoning (CBR) can meet this sustainability challenge, by supplementing a time series prediction model on a full-year-forecasting task. Using a dataset of three years of milk supply from Irish dairy farms (N=2,479), we produce accurate full-year forecasts for each individual farm, by augmenting that farm's data with data from nearest-neighboring farms, based on the similarity of their time series profiles (using Dynamic Time Warping). A study comparing four methods (Seasonal Naïve, LSTM, *Prophet*, $Prophet^{NN}$) showed that the method using CBR data-augmentation ($Prophet^{NN}$) outperformed the other evaluated methods. We also demonstrate the utility of CBR in providing farmers with novel *prefactual* explanations for forecasting that could help them to realize actions that could boost future milk yields and profitability.

Keywords: Smart Agriculture · Dairy Production · Time Series · Prefactual Explanation · CBR · Data Augmentation

1 Introduction

While SmartAg was originally predicated on delivering enhanced agriculture yields and productivity, increasingly it is becoming more about delivering a sustainable and efficient agriculture that minimally pollutes, delivers equivalent/better production levels from fewer inputs, aiming for a zero-carbon impact on the environment in accordance with the UN's sustainability goals and the promise of “AI for Good” [16, 42]. In the dairy sector, this challenge translates into producing the same or higher volumes of milk from fewer animals (e.g., genetically selecting cows that efficiently process what grass they eat), on pastures that make minimal use of artificial fertilizers (e.g., through using clover/grass mixes) and where carbon impacts are offset or balanced (e.g., using locally-grown grass rather than imported feed). Ultimately, meeting these challenges relies on understanding the relationships between a complex array of inputs (from animal genetics, to farm management to pollution measurement) and the volume of produce output by this sector, namely milk. Accurate long-term milk-supply forecasting plays a fundamental role in driving on-farm decision-making and processing capacity. Previous time series research has argued that *“combining of CBR with other approaches, seems promising and can improve the quality of forecasting”* [33]. CBR systems have also been successfully applied in oceanographic forecasting tasks [11]. More generally, in the agriculture domain, CBR solutions have enjoyed success in a variety of prediction tasks including grass growth prediction [24, 25, 40] and rangeland grasshopper infestation prediction [4, 15]. In this paper, we show the promise of AI techniques, in particular how time series analyses can be improved by data-augmentation techniques, using k-nearest neighbor methods from CBR [1], to accurately forecast long-term milk supply. In addition, we show how CBR can be used to generate *prefactual* explanations to help farmers realize actions that could be taken to boost milk yield in future years (Section 3), before concluding and discussing promising avenues for future research (Section 4).

1.1 Why Milk-Supply Forecasting is Important for Sustainability

In any given year, milk supply forecasting is a fundamental driver for the dairy sector. Dairy companies use their forecasts to establish pricing, contracts with farms, and the production requirements for their factories. As such, proper forecasting strongly influences on-farm management (e.g., in under/over production and manner of production), the consumption of resources in the sector (e.g., fertilizer use, tanker-transport use for milk collections) and the processing efficiency of factories (e.g., avoiding waste from surplus milk supplies) [36, 41]. Dairy processors can drive sustainability changes through accurate and precise forecasting. However, forecasting in this sector faces significant challenges. Milk-supply forecasts (i) must be made for 1000s of farms which differ from one another in their herd-profiles, the land farmed and their farm-management practices, (ii) have to be made in advance for the *full year*, for planning purposes, not incrementally as the year unfolds, (iii) can encounter disruption from climate-change and disease outbreaks (e.g., a hot summer can stop grass growth).

1.2 Predicting Milk Supply With Low-Resolution Data at Scale

Several models exist in the literature that can forecast milk supply accurately, but typically only for experimental farms, which have extensive and carefully-recorded data (e.g., on individual cows, farm management practices; see [45] for a review). In this single-farm prediction context, the most successful models are the surface-fitting model and the NARX (Nonlinear autoregressive model with exogenous inputs; which has a RMSE=75.5kg for a 365 day horizon [44]). These models use features such as Days-In-Milk (number of days a cow has been lactating) and the NCM (number of cows milked) in the herd. Some studies have used as many as 12 features, including genetics, feed, and grazing management information of the individual cows. However, such high-resolution data is rarely available for most commercial farms [45]. So, it is unclear how these forecasting methods can commercially scale to 1000s of farms.

In the present work, we attempt to forecast using low-resolution data where we have little information on individual cows and on-farm practices. In our time series model, we forecast for a given farm using the following minimal case-features: DATE (dd-mm-yyyy), WEEK (no. in year), MONTH (no. in year), HERD-DIM (Mean Days-in-Milk for the herd), CALVINGS (cumulative no. of calving's) to predict the target variable SUPPLY-QUANTITY (no. of litres in a given bulk-tanker collection). The DIM (Days-In-Milk) feature records the number of days a cow has been milking for and is important because a cow's milk yield varies over the season with a predictable profile (i.e., the so-called lactation curve that increases to Day-60 and then trails off (see [18]), The CALVINGS (cumulative no. of calving's) feature captures the number of offspring in the herd and is important because a cow only commences milking when it has given birth. Forecasting from these low-resolution inputs is quite hard because many key factors are missing (e.g., individual cow characteristics, amount and quality of feed used, calving management, etc.). Furthermore, the target variable, SUPPLY-QUANTITY, is not as informative as it could be, because it does not necessarily reflect one complete milking of the herd on a given day. It is a measure of the milk collected in a bulk tank on the farm (see Figure 1) by a tanker truck, and can reflect several milkings of the whole herd; specifically, as tanker collection times can vary between 1-3 days, a single collection could contain 1-5 complete milkings of the herd. Figure 2 shows the variability that can occur in these weekly milk-supply figures from one farm over the three-year period.

Fig. 1: An on-farm, bulk tank where cooled milk is stored, before being collected by a tanker truck by the milk processor.

1.3 How CBR Might Improve Milk Forecasting

Previous unpublished work on this milk forecasting problem, using low resolution data, has shown that deep learning methods can make better predictions than traditional auto-regressive methods. However, the absolute error-levels from these models were unacceptably high. The standard data-analytics approach to this problem uses historical milk-supply, time series data to predict on a farm-by-farm basis and then aggregates these predictions to get the forecast for all the farms supplying a single milk-processing company. Our conjecture was that the error found in previous models was due to one or other of three factors: (i) having insufficient data for a given farm, (ii) disruptive events in a farm's history that undermined generalizations across years (e.g., a once-off disease outbreak), (iii) changes in the tanker-collection schedule for a given farm. Hence, we hypothesised that, if the data for a single farm were supplemented with data from similar farms during the prediction step, then these sources of error might be reduced. This hypothesis invites a CBR solution, where we use k-NN to find a small number of nearest-neighbor farms based on some similarity metric (i.e., not necessarily spatially-proximate farms) and then use their data to augment

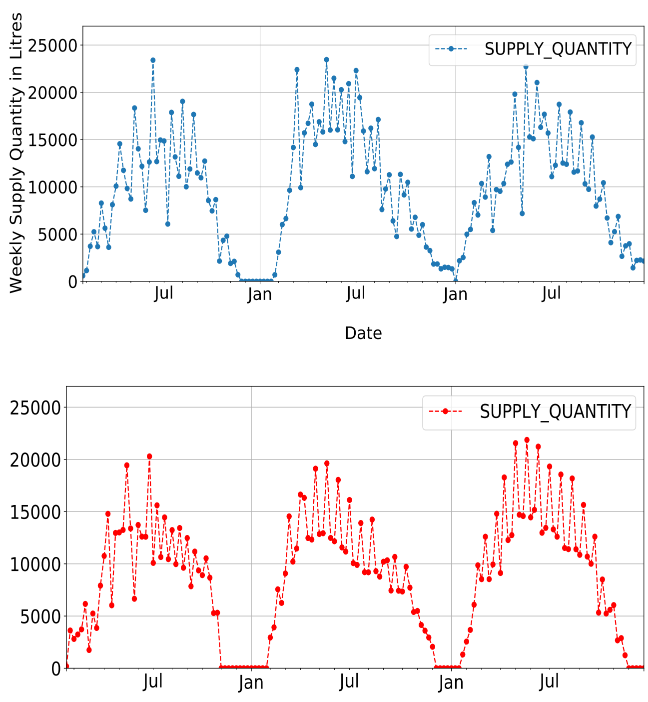

Fig. 2: Example of (a) 3-Year Milk Supply Profile of one farm, and (b) its nearest neighbor, retrieved using a dynamic time warping distance

the dataset for the time series predictions. A key question for this solution is how to determine the similarity between farms. Our solution bases this similarity on the multi-year profile of milk-supply between farms. One problem with this profile-similarity approach is that two farms may have very similar supply profiles but they may not align to one another in the time axis [2, 31], because one farm may have started milking slightly earlier/later in the year compared to the other. To solve this problem, we retrieve farms with a k-NN using a dynamic time warping (DTW) distance measure, which has previously been used in CBR systems for both classification and regression tasks [28]. DTW allows the matching of profiles irrespective of temporal-offsets that might occur in the year. Figure 2 shows one farm-profile over a three year period and its nearest neighbor when DTW is used.

2 A Study On Forecasting Milk Supply

Dataset. The data considered in this work covers dairy herds/farms (N=3,104) across 14 Irish counties for one dairy company (Glanbia) over four consecutive years. Cases describe the number of cows and calves in a herd at each milk collection and milk collected (target variable). On removing farms with missing data, the dataset had N=2,479 farms. Specific dates are anonymised throughout this paper on request from the industry partner furnishing the data.

Forecasting methods. Four different forecasting methods were used. First, as a simple benchmark [19], a Seasonal Naïve Forecasting method was used; it assumes that every week of the prediction-year's milk supply will be exactly the same as that of last year. While this is a popular benchmark in the forecasting space [19], it has surprisingly not been evaluated in previous milk supply prediction to the best of our knowledge (e.g., see review from [45]). Second, the *Long Short Term Memory (LSTM)* deep learning model was used as previously unpublished work found it to work best on this problem; LSTM stores sequences in long and short term states and then reuses them for prediction [17]. However, this has been shown to underperform basic statistical models on long-term forecasting and its processing overheads are an issue to be considered [29]. Third, the *Prophet* forecasting method, a generative additive model regularized with Bayesian techniques [38] was also used; it has been shown to work best for time series with strong seasonal effects and several seasons of historical data and in its simplest form is expressed as follows:

$$y(t) = g(t) + s(t) + \epsilon_t \quad (1)$$

Here the trend component, $g(t)$, automatically selects change points based on historic data by imposing a Laplacian prior. The seasonal component $s(t)$ is approximated with a Fourier series and a smoothing prior is imposed (see Equation 2). For yearly seasonality we set the regular period $P = 365.25$ and we set $N = 10$

allowing one to calculate the Fourier coefficients, a_n and b_n, as recommended by [38].

$$s(t) = \sum_{n=1}^{N}(a_n cos(\frac{2\pi nt}{P}) + b_n sin(\frac{2\pi nt}{P})) \tag{2}$$

The error term ϵ_t represents unusual changes that are not accommodated by the model but do contribute to the final forecast. An optional holiday term $h(t)$ can be included which accommodates strange but regular occurrences in a time series (e.g., a jump in activity around Christmas time when forecasting retail sales). Unlike many traditional forecasting techniques (e.g., ARIMA), *Prophet* automatically provides uncertainty estimates on its forecasts and does not require measurements to be regularly spaced in the time axis.

This technique has many useful *analyst-in-the-loop* characteristics that could be exploited in practical applications (e.g., specifying the maximum capacity of a farm, unusual events such as disease-outbreaks or seasonal-changes can be set in advance of prediction). Fourth and finally, our own $Prophet^{NN}$ was used; it extends the *Prophet* model by adding the k-nearest-neighboring farms (using DTW as a distance metric) to augment the training data prior to prediction (in the reported results k= 3 was used in this model). Specifically, the historic and predicted supply from the 3-NN herds are added as additional features in the model.

Evaluation metrics. Measures used included absolute error (AE) and mean absolute error (MAE), both measured in terms of litres-of-milk. We also considered MASE (Mean absolute scaled error), since it is the preferred evaluation technique in the forecasting literature, having many benefits over traditional measures [20]. When the MASE score is < 1 it means that the proposed method is outperforming the Seasonal Naïve forecast, whereas a MASE score > 1 means the opposite (i.e., smaller value better). Kullback-Leibler Divergence scores were used to compare model distributions to the ground truth.

Setup and evaluation. The first three years were used as training data and the final year as test data. Holdout strategies are preferred for real-world non-stationary time series data [6], where we want to maximize the ability of the models to learn seasonal effects from the full three years of training data. For each week of the test year we have a predicted value and the actual value. Forecasts were generated for each farm and then aggregated. For the LSTM implementation, the Adam optimizer was used with 100 epochs, a batch size of 4, and a learning rate of 0.001. The Keras API was used [7]. Dropout layers were implemented to prevent overfitting. The *Prophet* model was implemented using the month, week, herd average DIM, and cumulative calving number features, with milk supply quantity as the target variable. No additional hyperparameters were implemented. In $Prophet^{NN}$, k=3 for retrieving nearest neighbors for each farm when the forecasts were being made. All other $Prophet^{NN}$ parameters were identical to those used in *Prophet*.

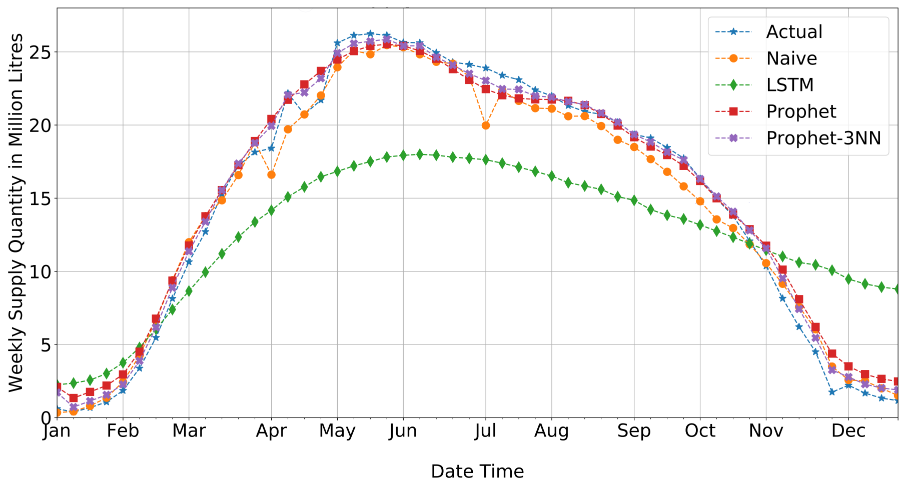

Fig. 3: A comparison of forecasts from the four methods relative to actual milk supply, as calculated on the final test year.

Table 1: Annual supply forecasts and error measures for each milk supply forecasting model. The ground truth value for actual production was 774.5M ltrs.

Evaluation	Unit	S-Naïve	LSTM	*Prophet*	*Prophet*^{NN}
Predicted Amount	Ltrs	755.8M	652.6M	796.1M	789.6M
Absolute Error	Ltrs	18.7M	129.1M	21.6M	15.1M
MAE (Weekly)	Ltrs	0.97M	4.8M	0.95M	0.58M
MASE	N/A	1	4.96	0.976	0.595

Results and discussion. The *Prophet*^{NN} model performs reliably better (MASE = 0.595) than both the original *Prophet* model (MASE = 0.976) and the Naïve forecast (MASE = 1.00; see Table 1). In terms of the absolute error over the whole year, the *Prophet*^{NN} prediction (789.6M ltrs) is closest to the actual value (774.5M ltrs), with a commercially acceptable error-margin (at 15.1M ltrs). It is 3.6M litres more accurate than the next best model (the Naïve model). This ordering of models by accuracy is mirrored by Kullback-Leibler Divergence scores, $D_{KL}(P||Q)$, comparing each model's distribution (Q) to the 2014 ground-truth distribution (P): *Prophet*^{NN} (0.003), Naïve (0.004), *Prophet* (0.008), LSTM (0.068). Notably, in milk supply prediction, small differences in error measures result in large differences in real-world, commercial outcomes (i.e., in millions of litres). In this respect, the Naïve model really does quite well, whereas the LSTM model is notably bad, confirming its under-performance in long-term forecasting [29]; however, it might improve with farm-by-farm hyperparameter tuning. All models tend to underpredict the peak milk-supply occurring in the summer months (see Fig.3). This is a common theme in milk supply forecasting [32,45]. However, $Prophet^{NN}$ performs best during this three-month period, (MAE = 0.38M). Both *Prophet* models overpredict milk supply at the start and end of the year, when supply is lowest, as farms begin and end milking for the year, when the Naïve model performs best (MAE = 0.37M). This result suggests that an ensemble approach may be useful in future work, where different models are used for different parts of the year. To summarize, the $Prophet^{NN}$ compares better than the other models evaluated in terms of MASE, MAE, AE and Kullback-Leibler Divergence score.

3 Providing Explainable Insights to Farmers

While providing forecasts on future milk supply from historical data can significantly enhance planning and resource allocation, the forecasts themselves do not provide any information on how a farm could increase its future output and boost profitability. One motivating reason for providing explanations to end-user farmers is to aid them in improving their current practice [23]. In the context of dairy farming this could translate into providing insights or recommendations that will help farmers to boost future milk output from the farm. Our focus in providing explanations to farmers is not on why a certain forecast is made based on historic data, but instead it is on realizing steps that could be taken to improve future yields.

3.1 Prefactual Explanations

A *prefactual explanation* describes a conditional (if-then) proposition about an, as yet not undertaken, action and the corresponding outcome that may (or may not) take place in the future [5,10]. While counterfactuals focus on past events, prefactuals center on the future and capture the idea of something that is not yet a fact, but could well become a fact [10].

In terms of goal planning, prefactual explanations can help individuals to determine how and whether a certain goal may be achieved in the future, and plan subsequent actions accordingly [10]. While counterfactual explanations have enjoyed success in providing explanations for past events (most commonly in classification systems [8,9,21,22,43]), prefactual explanations are relatively untapped, yet extremely promising for eXplainable AI (XAI) in forecasting scenarios. Perhaps the most relevant work here is in the area of goal-based recommendation, where CBR has been successfully applied to predict realistic new personal best race times for athletes and to recommend a suitable training plan to achieve their goals [12,37]. Similarly, in predicting forage-loss estimates due to grasshopper infestation, treatment recommendations have been provided that can help to minimize future economic loss [4]. These prefactual explanations could be used to provide insights to farmers that could help them to increase their milk output in future years. In the next sub-section, we describe a novel framework to formulate prefactual explanations, which is based on contrasting low-performing herds with high-performing *exemplar cases.*

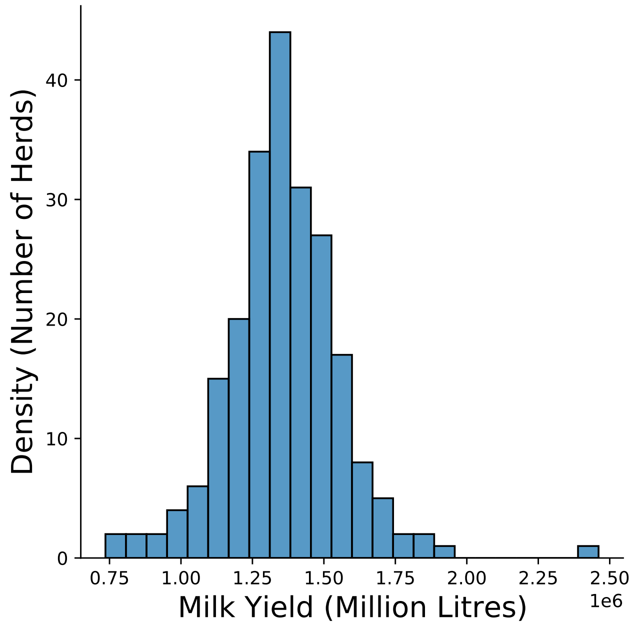

Fig. 4: The distribution of total milk output from medium size dairy farms (75-85 Cows, N=233 cases) over a four year period.

3.2 Using Prototypes in Prefactual Explanations

Prototypes are instances that are maximally representative of a class (typically retrieved using class centroids [30]) and have been used in several problem domains successfully to generate global explanations [13,26,30]. While there are no specific class labels in our problem, one observation from the case-base is that farms with a similar herd size (e.g., medium sized herds of $\approx$ 80 cows [39,41]) tend to vary greatly in terms of their milk output over the four-year period covered in the dataset (see Figure 4). So, for a fixed herd-size, there are both high performing herds with high milk yield, and low performing herds at the lower end of the distribution. By leveraging information from the high performing *exemplar cases* at the upper end of the distribution, lower-performing farms could modify their management practices with a view to enhancing future returns (e.g., an increased yield in the next year). More simply, prefactual explanations that use high-performing farms as exemplars should provide a basis for informing farmers on best practice. In the following sub-sections, we discuss the aspects of farm-management that could be the subject of these prefactual recommendations, with a view to identifying those that improve sustainability.

(I) Increasing herd size: As one might expect, there is a strong correlation between the number of animals on a farm and milk yield across the whole case-

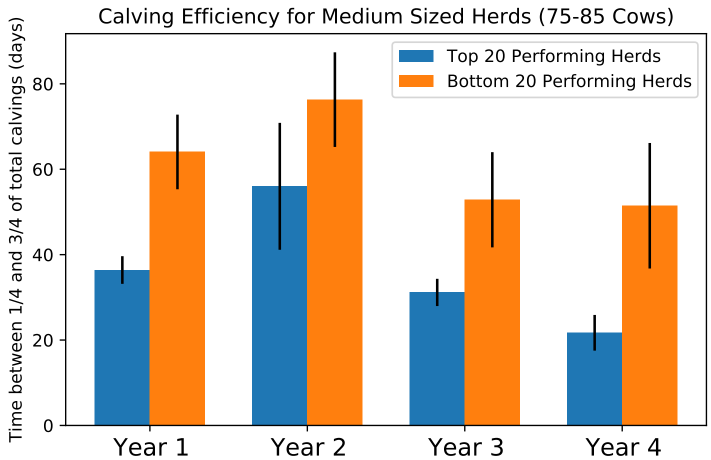

Fig. 5: Comparing both the most and the least efficient farms (at the tails of the distribution). The most efficient farms, in terms of milk output, calve much faster than the least efficient farms.

base (Pearson's $r = 0.947$). But, Figure 4 shows us that that many farms can increase their yield without increasing the size of their herd. There are several reasons why increasing herd size is not a good strategy for a farmer to take; (i) it requires new investment to buy more animals, (ii) it increases running costs of the herd (e.g., feed, fertilizer for grass, shelter, veterinary costs), estimated to be €1516 per animal on an Irish farm in 2022 [14, 35], (iii) it increases the likelihood of disease being introduced into a "closed" herd [35]. However, perhaps the biggest issue is that it is not a sustainable strategy. A prefactual could tell a farmer *"You could boost your yield by 10% next year if you increase your herd size by 10%"*, but such advice will also increase the carbon-costs of the farm and the potential for environmental damage and pollution (e.g., more cows equals more methane). As such, a more sustainable strategy would be to improve management practices without expanding the herd-size.

(II) Constraining calving: Dairy farmers using sustainable pasture-based systems where cows are mainly fed on grass in fields, are now strongly advised to control the timing of calving in their herds. Research has shown that constraining calving to a 6-week window in spring can greatly improve milk yield and herd fertility in pasture based-systems [35]. These benefits arise because after calving, cows begin milk production just as grass growth starts to peak; so, the animals ability to produce milk synchronises with the availability of sustainable feed (i.e., grass as opposed to bought-in, carbon-heavy supplementary feeds, such as soya or corn). Extended calving intervals result in the breakdown of the synchrony between feed supply and feed demand and often result in reduced fertility, over-fat cows and an increased likelihood of disease in the herd [35]. All of these factors negatively impact milk yield and often lead to increases in artificial insemination, veterinary and hormonal costs [3, 35]. Indeed, this proposed relationship between calving-periods and milk yield is supported by analyses of our dataset. We estimated on-farm calving efficiency by analysing the mean amount of time taken between the calving of the 1/4 and 3/4 of calves born on farms across the four year period. Using a fixed herd size for (i) the top 20 performing *exemplar* herds, and (ii) the worst 20 performing herds (see e.g., Figure 5), it was found that the best or *exemplary* herds calve significantly faster than herds with lower milk yields (one sided t-test, $p < 0.001$). Therefore, one useful prefactual explanation for farmers hoping to boost subsequent yield could be of the form; *"If you constrain calving to a shorter period than last year, your future milk supply is likely to increase"*.

(III) Optimizing the supply profile: To gain insights into the optimal supply profile for a herd, we compare the prototypical supply profile for both high-performing herds and herds with the lowest supply for a fixed herd-size. Prototypical herds are created using centroids from: (i) k-medoids clustering with a DTW distance measure as it retrieves realistic instances that are already part of the case-base [30], and (ii) k-means clustering with dynamic barycenter averaging as suggested by [34]. This analysis showed that *exemplary* herds have a much shorter drying-off period compared to the weaker performing herds (See Figure 6). This "drying-off" period (which usually lasts 6-8 weeks) typically occurs at the end of the year, when cows are not milked, to allow them to recover health-wise before the next calving season [27]. Although this period plays a critical role in rejuvenating an animal's health, too long a dry period can result in over-fat cows and reduced fertility. So, an extended dry period can reduce milk yield on a given farm. While research suggests that drying off decisions should be made on an individual cow basis [27], we found that lower performing herds tend to have over-extended periods of no-milkings, hurting their yields. The prototypical supply curves for *exemplar* herds, unlike low-yielding herds, quickly grow to reach their peak yield in early summer and slowly taper off towards the dry period in late November/early-December. Indeed, their milk-supply curves track the grass growth curves observed on farms during these periods (see Kenny et al., [25]), where grass growth typically peaks in early May before tapering off into Autumn. Therefore, one useful prefactual explanation for farmers hoping to boost subsequent yield could be of the form; *"If you slightly reduce the drying off period, (i) your future milk supply is likely to increase as you will have longer milking periods, (ii) you could be less reliant on purchasing feed as the supply curve is more likely to track the grass growth curve."*.

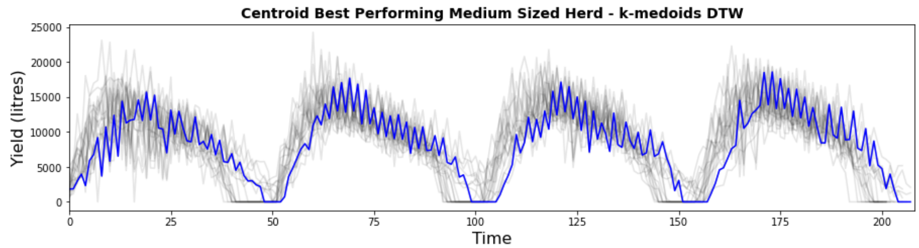

(a)

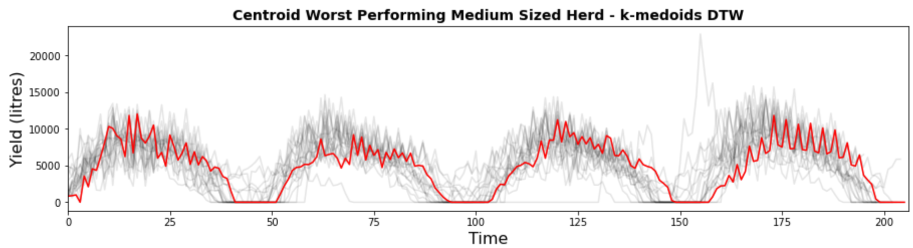

(b)

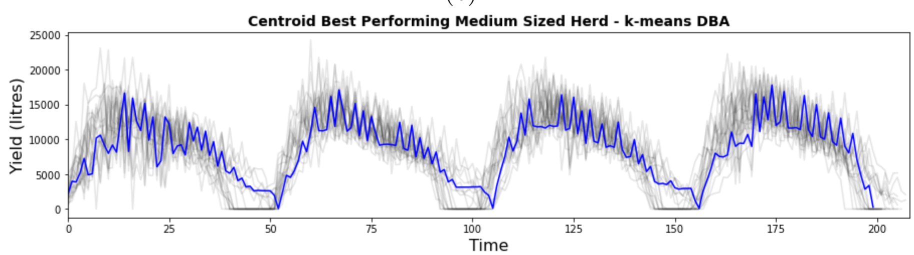

(c)

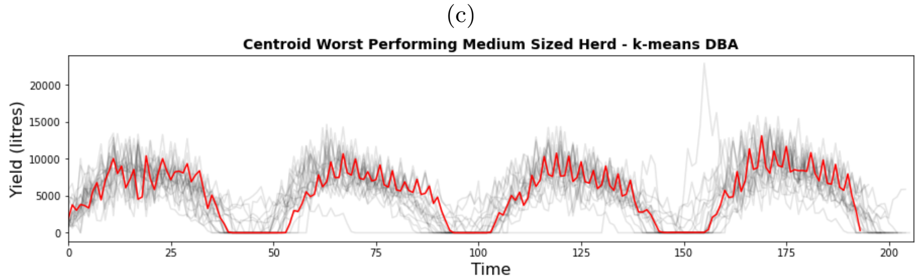

(d)

Fig. 6: Prototypical milk yield supply profiles over a four year period for high (in blue) and low (in red) performing medium-sized herds according to prototypes retrieved using the centroid from: (i) k-medoids DTW shown in (a) & (b), and (ii) k-means DBA in (c) & (d).

4 Conclusion and Future Directions

Facing the challenge of developing a precision agriculture that can support sustainability, we have demonstrated a promising CBR data-augmentation technique, using nearest neighboring farms with similar production profiles, which makes acceptably accurate long-term forecasts for milk supply based on low resolution data. As milk-supply forecasting drives significant aspects of the agricultural sector, better forecasting can play a critical role in reducing resource waste from farm to factory and in budgeting-on farm. An immediate avenue for future work is to investigate hyperparamater tuning on a validation set in an attempt to boost predictive performance. We also explored the utility of CBR in providing novel goal-oriented *prefactual* explanations to farmers to help them realize actions that could boost milk yield in future years. A core novelty of this work is in generating *prefactual* explanations through leveraging high performing *exemplar* cases, and in future work we would like to develop alternative algorithmic approaches to generate such explanations and to explore their utility in domains beyond agriculture. All of the current work, demonstrates how CBR and AI has the potential to improve farm-management practices in ways that deliver a more efficient, less polluting and more sustainable agriculture.

Acknowledgements. This publication has emanated from research conducted with the financial support of (i) Science Foundation Ireland (SFI) to the Insight Centre for Data Analytics under Grant Number 12/RC/2289_P2 and (ii) SFI and the Department of Agriculture, Food and Marine on behalf of the Government of Ireland under Grant Number 16/RC/3835 (VistaMilk).

References

1. Aha, D.W., Kibler, D., Albert, M.K.: Instance-based learning algorithms. Machine learning **6**(1), 37–66 (1991)
2. Berndt, D.J., Clifford, J.: Using dynamic time warping to find patterns in time series. In: KDD workshop. vol. 10, pp. 359–370. Seattle, WA, USA: (1994)
3. Boichard, D.: Estimation of the economic value of conception rate in dairy cattle. Livestock Production Science **24**(3), 187–204 (1990)
4. Branting, K., Hastings, J.D., Lockwood, J.A.: Carma: A case-based range management advisor. In: IAAI. pp. 3–10 (2001)
5. Byrne, R.M., Egan, S.M.: Counterfactual and prefactual conditionals. Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale **58**(2), 113 (2004)
6. Cerqueira, V., Torgo, L., Mozetič, I.: Evaluating time series forecasting models: An empirical study on performance estimation methods. Machine Learning **109**(11), 1997–2028 (2020)
7. Chollet, F., et al.: Keras (2015), `https://github.com/fchollet/keras`
8. Dai, X., Keane, M.T., Shalloo, L., Ruelle, E., Byrne, R.M.: Counterfactual explanations for prediction and diagnosis in xai. In: To appear in AIES'22 (2022)

9. Delaney, E., Greene, D., Keane, M.T.: Instance-based counterfactual explanations for time series classification. In: International Conference on Case-Based Reasoning. pp. 32–47. Springer (2021)
10. Epstude, K., Scholl, A., Roese, N.J.: Prefactual thoughts: Mental simulations about what might happen. Review of General Psychology **20**(1), 48–56 (2016)
11. Fdez-Riverola, F., Corchado, J.M.: Fsfrt: Forecasting system for red tides. Applied Intelligence **21**(3), 251–264 (2004)
12. Feely, C., Caulfield, B., Lawlor, A., Smyth, B.: Using case-based reasoning to predict marathon performance and recommend tailored training plans. In: International Conference on Case-Based Reasoning. pp. 67–81. Springer (2020)
13. Gee, A.H., Garcia-Olano, D., Ghosh, J., Paydarfar, D.: Explaining deep classification of time-series data with learned prototypes. CEUR Workshop Proceedings **2429**, 15–22 (2019)
14. Gowing, P., Dunne, J.: Cost control for 2022 (2022), `https://www.teagasc.ie/news--events/daily/dairy/cost-control-for-2022.php`
15. Hastings, J., Branting, K., Lockwood, J.: Carma: A case-based rangeland management adviser. AI Magazine **23**(2), 49–49 (2002)
16. Herweijer, C., Combes, B., Gillham, J.: How ai can enable a sustainable future. PwC report (2018)
17. Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation **9**(8), 1735–1780 (1997)
18. Horan, B., Dillon, P., Berry, D., O'Connor, P., Rath, M.: The effect of strain of holstein–friesian, feeding system and parity on lactation curves characteristics of spring-calving dairy cows. Livestock Production Science **95**(3), 231–241 (2005)
19. Hyndman, R.J., Athanasopoulos, G.: Forecasting: principles and practice. OTexts (2018)
20. Hyndman, R.J., Koehler, A.B.: Another look at measures of forecast accuracy. International journal of forecasting **22**(4), 679–688 (2006)
21. Keane, M.T., Kenny, E.M., Delaney, E., Smyth, B.: If only we had better counterfactual explanations: Five key deficits to rectify in the evaluation of counterfactual xai techniques. In: IJCAI-21 (2021)
22. Keane, M.T., Smyth, B.: Good counterfactuals and where to find them: A case-based technique for generating counterfactuals for explainable ai (xai). In: Proc. ICCBR'20. pp. 163–178. Springer (2020)
23. Kenny, E.M., Ford, C., Quinn, M., Keane, M.T.: Explaining black-box classifiers using post-hoc explanations-by-example: The effect of explanations and error-rates in xai user studies. Artificial Intelligence **294**, 103459 (2021)
24. Kenny, E.M., Ruelle, E., Geoghegan, A., Shalloo, L., O'Leary, M., O'Donovan, M., Temraz, M., Keane, M.T.: Bayesian case-exclusion and personalized explanations for sustainable dairy farming. In: IJCAI. pp. 4740–4744 (2021)
25. Kenny, E.M., Ruelle, E., Geoghegan, A., Shalloo, L., O'Leary, M., O'Donovan, M., Keane, M.T.: Predicting grass growth for sustainable dairy farming: A cbr system using bayesian case-exclusion and post-hoc, personalized explanation-by-example (xai). In: International Conference on Case-Based Reasoning. pp. 172–187. Springer (2019)
26. Kim, B., Khanna, R., Koyejo, O.O.: Examples are not enough, learn to criticize! criticism for interpretability. In: Advances in Neural Information Processing Systems. pp. 2280–2288 (2016)
27. Kok, A., Chen, J., Kemp, B., Van Knegsel, A.: Dry period length in dairy cows and consequences for metabolism and welfare and customised management strategies. animal **13**(S1), s42–s51 (2019)

28. Mahato, V., O'Reilly, M., Cunningham, P.: A comparison of k-nn methods for time series classification and regression. In: AICS. pp. 102–113 (2018)
29. Makridakis, S., Spiliotis, E., Assimakopoulos, V.: Statistical and machine learning forecasting methods: Concerns and ways forward. PloS one **13**(3), e0194889 (2018)
30. Molnar, C.: Interpretable machine learning. Lulu.com (2020)
31. Mueen, A., Keogh, E.: Extracting optimal performance from dynamic time warping. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp. 2129–2130 (2016)
32. Murphy, M.D., O'Mahony, M.J., Shalloo, L., French, P., Upton, J.: Comparison of modelling techniques for milk-production forecasting. Journal of dairy science **97**(6), 3352–3363 (2014)
33. Nakhaeizadeh, G.: Learning prediction of time series. a theoretical and empirical comparison of cbr with some other approaches. In: European Workshop on Case-Based Reasoning. pp. 65–76. Springer (1993)
34. Petitjean, F., Forestier, G., Webb, G.I., Nicholson, A.E., Chen, Y., Keogh, E.: Dynamic time warping averaging of time series allows faster and more accurate classification. In: 2014 IEEE international conference on data mining. pp. 470–479. IEEE (2014)
35. Shalloo, L., Cromie, A., McHugh, N.: Effect of fertility on the economics of pasture-based dairy systems. Animal **8**(s1), 222–231 (2014)
36. Shalloo, L., Creighton, P., O'Donovan, M.: The economics of reseeding on a dairy farm. Irish Journal of Agricultural and Food Research pp. 113–122 (2011)
37. Smyth, B., Cunningham, P.: A novel recommender system for helping marathoners to achieve a new personal-best. In: Proceedings of the eleventh ACM conference on recommender systems. pp. 116–120 (2017)
38. Taylor, S.J., Letham, B.: Forecasting at scale. The American Statistician **72**(1), 37–45 (2018)
39. Teagasc: Teagasc national farm survey 2020 - dairy enterprise fact-sheet (2020), `https://www.teagasc.ie/media/website/publications/2021/NFS_Dairy_Factsheet2020.pdf`
40. Temraz, M., Kenny, E.M., Ruelle, E., Shalloo, L., Smyth, B., Keane, M.T.: Handling climate change using counterfactuals: Using counterfactuals in data augmentation to predict crop growth in an uncertain climate future. In: International Conference on Case-Based Reasoning. pp. 216–231. Springer (2021)
41. Upton, J., Murphy, M., De Boer, I., Koerkamp, P.G., Berentsen, P., Shalloo, L.: Investment appraisal of technology innovations on dairy farm electricity consumption. Journal of dairy science **98**(2), 898–909 (2015)
42. Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S.D., Tegmark, M., Fuso Nerini, F.: The role of artificial intelligence in achieving the sustainable development goals. Nature communications **11**(1), 1–10 (2020)
43. Wiratunga, N., Wijekoon, A., Nkisi-Orji, I., Martin, K., Palihawadana, C., Corsar, D.: Actionable feature discovery in counterfactuals using feature relevance explainers. Third Workshop on XCBR: Case-based reasoning for the explanation of intelligent systems. (2021)
44. Zhang, F., Murphy, M.D., Shalloo, L., Ruelle, E., Upton, J.: An automatic model configuration and optimization system for milk production forecasting. Computers and electronics in agriculture **128**, 100–111 (2016)
45. Zhang, F., Shine, P., Upton, J., Shaloo, L., Murphy, M.D.: A review of milk production forecasting models: Past & future methods (2020)