


HYPO-CHEAT's aggregated weekly visualisations of risk reduce real world hypoglycaemia

Digital Health
Volume 8: 1–22
© The Author(s) 2022
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/20552076221129712
journals.sagepub.com/home/dhj



Chris Worth^{1,2} , Paul W Nutter¹, Mark J Dunne³,
Maria Salomon-Estebanez², Indraneel Banerjee^{2,3}
and Simon Harper¹

Abstract

Background: Children with congenital hyperinsulinism (CHI) are at constant risk of hypoglycaemia with the attendant risk of brain injury. Current hypoglycaemia prevention methods centre on the prediction of a continuous glucose variable using machine learning (ML) processing of continuous glucose monitoring (CGM). This approach ignores repetitive and predictable behavioural factors and is dependent upon ongoing CGM. Thus, there has been very limited success in reducing real-world hypoglycaemia with a ML approach in any condition.

Objectives: We describe the development of HYPO-CHEAT (**HY**poglycaemia-**P**revention-thr**O**ugh-CGM-**HE**atmap-**T**echnology), which is designed to overcome these limitations by describing weekly hypoglycaemia risk. We tested HYPO-CHEAT in a real-world setting to evaluate change in hypoglycaemia.

Methods: HYPO-CHEAT aggregates individual CGM data to identify weekly hypoglycaemia patterns. These are visualised via a hypoglycaemia heatmap along with actionable interpretations and targets. The algorithm is iterative and reacts to anticipated changing patterns of hypoglycaemia. HYPO-CHEAT was compared with Dexcom Clarity's pattern identification and Facebook Prophet's forecasting algorithm using data from 10 children with CHI using CGM for 12 weeks. HYPO-CHEAT's efficacy was assessed via change in time below range (TBR).

Results: HYPO-CHEAT identified hypoglycaemia patterns in all patients. Dexcom Clarity identified no patterns. Predictions from Facebook Prophet were inconsistent and difficult to interpret. Importantly, the patterns identified by HYPO-CHEAT matched the lived experience of all patients, generating new and actionable understanding of the cause of hypos. This facilitated patients to significantly reduce their time in hypoglycaemia from 7.1% to 5.4% even when real-time CGM data was removed.

Conclusions: HYPO-CHEAT's personalised hypoglycaemia heatmaps reduced total and targeted TBR even when CGM was reblinded. HYPO-CHEAT offers a highly effective and immediately available personalised approach to prevent hypoglycaemia and empower patients to self-care.

Keywords

Hypoglycaemia, prevention, human-computer interaction, data aggregation, behaviour change, continuous glucose monitoring

Submission date: 28 March 2022; Acceptance date: 13 October 2021

¹Department of Computer Science, University of Manchester, Manchester, UK

²Department of Paediatric Endocrinology, Royal Manchester Children's Hospital, Manchester, UK

³Faculty of Biology, Medicine and Health, University of Manchester, Manchester, UK

Corresponding author:

Chris Worth, Department of Paediatric Endocrinology, Royal Manchester Children's Hospital, Oxford Road, Manchester, M13 9WL, UK.
Email: chris.worth@mft.nhs.uk



Introduction

Congenital hyperinsulinism (CHI) is a disorder of unregulated and unpredictable insulin secretion and is the commonest cause of severe and unpredictable hypoglycaemia in children. Hypoglycaemia (hypo) refers to blood glucose levels below the lower limit of normal and, in children with CHI, a pragmatic cut-off of 3.5mmol/l is frequently used.^{1,2} Recurrent hypoglycaemia in early childhood from any cause can lead to severe long-term impairments such as epilepsy, hemiparesis and mental retardation.³ In CHI, rates range from 15% to almost 50%.^{4,5} Hypoglycaemia has acute clinical impact and delayed consequential impact on quality of life and the healthcare economy.^{6,7}

Monitoring and review of glucose levels

The current method of hypoglycaemia monitoring, detection and prevention for most patients worldwide is intermittent fingerprick testing. However, the lack of trend information provides little opportunity to truly predict and prevent episodes and thus there is an opportunity for missed hypoglycaemia between tests. As a consequence of this, continuous glucose monitoring (CGM) has emerged as a popular option in recent years. CGM devices monitor interstitial glucose levels and report a value to a receiver device every five minutes. While CGM does improve time in range (TIR) for those with type 1 diabetes mellitus (T1DM),⁸ there is no evidence of this in patients with CHI and its ability to predict the onset of impending hypoglycaemia is limited.⁹ Currently available methods for the users to retrospectively review CGM data (e.g. Dexcom Clarity) present large volumes of difficult to understand information with minimal interpretation and no clear actions to reduce hypos. While little attention has been paid to improving the interpretability and actionability of retrospective CGM review, there has been great interest in using modern algorithmics to acutely predict future glucose values and thus the onset of impending hypoglycaemia.

Prediction of glucose (and thus hypoglycaemia)

The majority of this research has focused on using machine learning (ML) and continuous CGM (in contrast to intermittent CGM: short periods of use to understand individual patterns and level of control) to predict a continuous glucose variable with a short prediction horizon of 30 minutes.¹⁰ This approach has multiple problems and so far has not demonstrated a reduction of hypoglycaemia in the real world. The limitations of this approach are discussed in detail in the Related work' section but are briefly summarised here. These algorithms invariably focus on the use of physiological measures to predict a continuous variable,

ignoring the fact that patterns of hypoglycaemia are dependent on the hour of day and day of week¹¹⁻¹³ and thus are influenced by repetitive behaviours rather than physiology. Hypoglycaemia risk is assumed to be consistent across days and thus all days are treated equally despite the differences created by routines such as school, work, hobbies and exercise. These algorithms offer a prediction of future glucose values up to 30 minutes ahead¹⁰ but, given device lag and the time required to find and administer treatment (glucose), this information is rarely actionable as this is often too short a time to act to prevent a rapidly developing hypoglycaemia in children with CHI. Furthermore, no underlying reasons for the hypoglycaemia are provided and so no lessons are learnt for the future. Finally, these algorithms all require continuous CGM to function and are thus expensive and not available to the majority of patients worldwide.¹⁴ Access to real-time CGM data from device manufacturers is possible in certain cases but certainly far from straightforward for the majority of research teams outside commercial studies. This significantly restricts the ability to test algorithms in the real world¹⁵ and the vast majority are only evaluated in silico with no assessment of their ability to reduce hypoglycaemia for real patients in free-living conditions.

Our approach

With no readily available way to easily interpret CGM data in an actionable format, and ML-driven glucose forecasting approaches to prediction severely limited, a new and complementary approach was required. Based on formative work demonstrating predictable patterns in hypoglycaemia risk in patients with CHI^{11,13} we hypothesised that a novel approach that aggregates data to calculate discrete weekly hypoglycaemia risk and concisely presents this visually to patients and with limited but defined actions would be of great practical use (Figure 1).

This was the basis for the development of HYPO-CHEAT (**HY**poplycaemia-**P**revention-thr**O**ugh-CGM-**HE**atmap-**T**echnology), which aggregates several weeks of individual CGM data to generate the hourly risk of hypoglycaemia for each separate day in a week. This is visualised as a personalised hypoglycaemia risk heatmap along with text-based interpretations and actionable suggestions for reflection. This user-friendly visualisation is shown to patients who can use this data to proactively prevent hypoglycaemia in the upcoming weeks without the need for ongoing CGM. This iterative and proactive approach stands HYPO-CHEAT apart from other hypoglycaemia prevention algorithms and available CGM review software which are frequently reactive in their provision of non-actionable data and thus fail to look forward to proactively prevent hypos. HYPO-CHEAT is designed to be either used alone when access to CGM is limited, or, in resource-rich environments, as a complementary approach

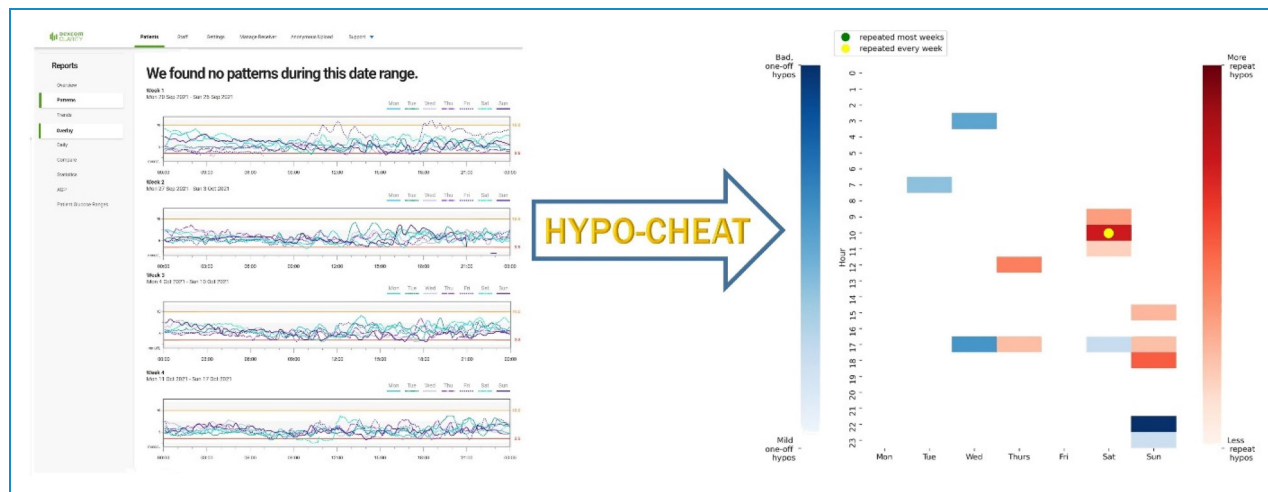


Figure 1. Graphical representation of HYPO-CHEAT's approach. On the left is the Dexcom Clarity output for 4 weeks of CGM data from one user. No patterns have been identified and there are no suggestions to aid improvement or actionable data. On the right is the HYPO-CHEAT output for this same period, demonstrating a succinct and clear visualisation of aggregated data, presented in a way to facilitate reflection and prompt simple action to improve outcomes. This visualisation would be accompanied by a suggestion to focus on Saturday morning and Sunday afternoon.

CGM: continuous glucose monitoring; HYPO-CHEAT: **HY**poglycaemia-**P**revention-thr**O**ugh-**CGM**-**H**eatmap-**T**echnology.

to recursive algorithms offering acute hypoglycaemia prediction.

We performed a pilot evaluation of HYPO-CHEAT in 10 patients with CHI who were using CGM for 12 weeks in free-living conditions. HYPO-CHEAT identified weekly patterns of hypoglycaemia and, when asked to reflect on hotspots, patients were able to identify repetitive behaviours associated with hypoglycaemia events. Neither Dexcom Clarity nor Facebook Prophet reliably identified these patterns. Despite the removal of CGM data from patients (by blinding the device), use of HYPO-CHEAT significantly reduced the time patients spent hypoglycaemic (time below range (TBR)).

Using a novel aggregation of data on weekly hypoglycaemia risk and the presentation of this in a concise and actionable format, our approach can reduce TBR for patients with CHI in addition to obviating the need for ongoing CGM and thus reducing cost (albeit still requiring a period of high-cost CGM) and patient burden.

Related work

The focus of this work intersects with other methods of hypoglycaemia prevention which primarily fall into one of three categories: Glucose forecasting; Decision Support Systems (DSSs) and Provision and analysis of CGM data. The artificial pancreas is an effective area within T1DM for hypoglycaemia prevention but is not relevant to our patient population where the primary focus is hypoglycaemia and thus we do not discuss it. Structured education is an effective technique to prevent hypoglycaemia but an in-depth discussion of this is

of limited use when discussing algorithmic approaches and is therefore not provided.

ML approaches

ML has become a dominant force in hypoglycaemia prediction and prevention in recent years¹⁶ and is now the primary tool used in Glucose forecasting and DSSs.

Glucose forecasting. Glucose forecasting algorithms aim to prevent hypoglycaemia by continuously predicting a glucose value within the prediction horizon (normally 30 minutes) and thus alerting the patient to any values that may fall below a threshold. Simple time series models have shown good in silico performance¹⁷ and some data would suggest that these algorithms perform as well as those driven by ML.¹⁸ However, ML now forms the bulk of this work and these algorithms come in one of three formats: physiological; data-driven; and hybrid.¹⁰ Data-driven and hybrid are of interest to the computer science community, being driven by ML in a 'black-box' approach to variable prediction. Various individual groups have had some success in using ML to predict a future glucose value in silico^{19,20} and, within the artificial pancreas environment, even in real subjects (under strict laboratory conditions with exercise and insulin as additional algorithmic inputs).²¹ However, a meta-analysis reported an overall sensitivity and specificity of 79% and 80% for prediction of hypoglycaemia in silico from all ML approaches.²² It was concluded in both this meta-analysis and a further systematic review¹⁶ that current ML

algorithms are currently not sufficiently accurate to be effective in a clinical setting.

Unfortunately, the vast majority of studies are conducted *in silico* rather than clinically²³ and this is likely to enhance reported accuracy over real-world testing due to the lack of real-world variables confounding glycaemic predictions.²⁴ Most importantly, the black box approach of ML precludes any reflection on the hypoglycaemia events²⁵ and thus prevents the development of a better understanding of the patterns or causes of hypoglycaemia upon which the patient could act. These algorithms also ignore the human in the loop and assume that mere provision of data will prevent hypoglycaemia despite good evidence demonstrating that poor self-management (rather than a lack of information) is the commonest cause of hypoglycaemia²⁶ and mere provision of data does not prevent events.¹¹

DSSs. DSSs improve glucose forecasting by recognising the human in the loop and aim to facilitate decision making using algorithmics based on various inputs and predictions. Various DSSs have reported successful outcomes but either require hugely complex inputs from users²⁷ and/or are only evaluated *in silico*.²⁸ Both Skrovseth et al.²⁹ and Liu et al.³⁰ showed real-world reductions in hypoglycaemia with their DSSs, but showed no additional benefit from the addition of ML algorithmics, aptly demonstrating that what really helped to reduce hypoglycaemia was feedback and reflection.

DSSs are plagued by the same issues as glucose forecasting algorithms, namely poor interpretability; *in silico* testing; assumptions that all recommendations are followed without focusing on how to achieve this; and incomplete feature sets. Thus, Tyler and Jacobs conclude in their review of ML in DSSs, that ‘it has not yet been shown that a DSS can improve TIR in human studies’.²⁴

Provision and analysis of CGM

Rather than use algorithmics to try and predict hypoglycaemia, it is possible to simply provide patients with a CGM device and assume that the availability of live data will translate into reduced hypoglycaemia. While there is evidence that continuous CGM improves TIR for those with T1DM on a large scale,^{8,31} evidence is sadly lacking for those with non-hypoglycaemia disorders such as CHI.⁹ From our prior work, we know that simple provision of CGM is not sufficient to prevent hypoglycaemia,¹¹ likely due to the multiple barriers patients face: lag time³², alarm fatigue³³ and poor accuracy (see the ‘Accuracy of CGM’ section).

Accuracy of CGM. Accuracy is a particular problem for acute detection and prevention of hypoglycaemia and is significantly worse for those with non-diabetes hypoglycaemia.^{34,35} Studies conducted exclusively on patients with CHI show mean absolute relative difference (MARD) of

17.9%³⁶ and 17.5%³⁷ and a hypoglycaemia (<3.9mmol/l) sensitivity of only 43%.³⁷ Our own analysis has demonstrated a MARD of 19.3% and hypoglycaemia (glucose <3.5mmol/l) sensitivity of 44% with the latest available device (Dexcom G6).³⁸ This suggests that if patients rely solely on acute hypoglycaemia detection using CGM they risk more than half of hypoglycaemia episodes going undetected by the device. It is clear that this reactive system is not sufficient to keep patients safe and a proactive approach is required. The use of CGM as a phenotyping tool is less sensitive to suboptimal accuracy as the stress is placed on patterns rather than point accuracy. Since accuracy does not vary by time of day, patterns of hypoglycaemia can be relied upon to be a reasonable representation of the true picture with the added advantage of offering actionable data well ahead of the likely hypoglycaemia.

Review of CGM data. With extended use of CGM, it might be expected that patients would come to understand their patterns of hypoglycaemia and act to avoid these. Unfortunately, the currently available methods for reviewing CGM data are cumbersome, data heavy and do not synthesise data into easy-to-interpret, actionable outputs from which patients could inform behaviour change to prevent hypos. Additionally, CGM is often prohibitively expensive and rarely funded by health services for patients with CHI.¹⁴ Thus, the provision of continuous CGM in CHI is unfortunately beyond the reach of many patients.

Our approach

We have outlined the problems with the current approaches to hypoglycaemia prevention. Below, we describe our novel approach in patients with CHI and examine how this differs from the work discussed above.

Experimental prototype (HYPO-CHEAT)

The fundamental problem with current approaches is that they are either reactive (ML-based prediction) or provide an overwhelming amount of non-actionable data (CGM review). They are blind to repetitive and predictable behaviour and assume that the provision of information will automatically result in an improvement in outcomes. To ensure HYPO-CHEAT does not stumble into these pitfalls, we outline our approach as follows:

- **Aggregation of data into discrete hourly buckets** rather than the prediction of a continuous variable. This allows for future prediction of hypoglycaemia and is proactive rather than reactive.
- Recognition of **patterns of hypoglycaemia that can be associated with behaviours** and thus are amenable to reflection on their causes and provide an element of education.

- Visualisation of a select amount of clear data which is easy to interpret and **can be easily actioned** by patients; this is in contrast to the large volumes of unedited data with no clearly associated actions provided by current CGM review websites.
- **An iterative system** which not only adapts to changing patterns of behaviour and hypoglycaemia but expects and looks for this change following suggestions.
- A system that is informed by CGM but **does not rely on continuous CGM for long periods**, thus reducing cost and patient burden as well as increasing the ease of real-world evaluation.

Generation of the best possible dataset. Due to the imperfect accuracy of CGM, patients often show short fluctuations (one or two values representing 5–10 minutes) below a hypoglycaemia threshold. These normally do not represent true hypos. In order to increase the proportion of hypoglycaemia analysed that is genuine, HYPO-CHEAT defines hypos as per the American Diabetes Association (ADA) criteria³⁹: three consecutive CGM values below a threshold (>15 minutes) to commence a hypo and three above a threshold (>15 minutes) to end the hypo. Once these criteria are met, then the first below-threshold value is taken as the start time and the first above-threshold value as the end time.

To further increase the proportion of hypos that are genuine, all fingerprick glucometer values (higher level of accuracy) are compared with CGM hypos and if glucometer values lie within the time period of a CGM reported hypo but are above the threshold, then this episode is removed.

Aggregation of data to detail time in hypoglycaemia. As detailed above, there are multiple problems with trying to prevent hypoglycaemia using ML and CGM to predict a continuous variable based on physiological measures. HYPO-CHEAT overcomes each of these issues by approaching the problem from a novel direction.

HYPO-CHEAT is designed to solely assess for hypoglycaemia. Glucose values above the set threshold can add unnecessary complexity and are therefore ignored. HYPO-CHEAT does not attempt to predict a glucose value at any given time but rather aggregates data to detail the frequency of hypos by every hour of every day. After the initial phenotyping period of several weeks, HYPO-CHEAT has generated a profile of repeating hypo risk (by each hour of each day) for each patient which is likely to reflect individual behaviour rather than simple physiology. This digital phenotype also highlights periods that contain hypos every (or almost every) week and thus, given the repetitive nature of human behaviour, can be used to predict, and hopefully prevent, the following week's hypos.

Formulation of an appropriate window for display and iterative analysis. The next step was to understand the repeating window within which the aggregated data sits. Given the social construct of the 7-day week and the organisation of most people's lives around this idea, it was felt that seven days was likely to offer a useful repeating window for the assessment of hypoglycaemia risk. There is support in the literature for this hypothesis; Dave et al. reported that day of week was a strong predictor of glucose values.¹² To ensure that this was the case for our data, we used Facebook Prophet time series forecasting to assess for patterns in CGM data from 10 patients.

Facebook Prophet is designed to predict a continuous variable such as glucose value. To avoid a potential negating effect of high glucose values on low glucose values, we converted the input to percentage time hypoglycaemic/hour. This provided a more useful prediction than an average glucose value. The averaging effect of this algorithm results in a predicted negative per cent time hypoglycaemic for some periods and is thus of limited use for predicting future periods of risk. However, its analytical power does provide good pattern recognition and thus allows us to identify repeating patterns. For all patients, this showed a daily and weekly pattern to the data and confirmed our hypothesis (Figure 2). Thus, HYPO-CHEAT uses a weekly window for visualisation and iterative analysis of change to the aggregated data described in the 'Aggregation of data to detail time in hypoglycaemia' section.

This approach has multiple benefits. HYPO-CHEAT is fundamentally based on the idea that humans are creatures of habit and will engage in regular activities that increase hypoglycaemia risk at certain time points in the week (e.g. exercise, prolonged fasting and repeating meal habits).^{40,41} These may not be at the same time each day (e.g. many people's routines will be different on a Tuesday to a Saturday) and thus are missed if one does not consider each day as a different entity. HYPO-CHEAT considers all days as individual periods and thus can spot patterns that occur at specific times on specific days that would be missed otherwise. Provision of a risk profile for the entire upcoming week allows patients to plan sufficiently in advance to prevent their hypo and also allows for reflection on the causes of these repeated hypos. The final advantage is one of cost and patient burden: after a period of CGM sufficiently long to generate a reliable phenotype, the CGM can be discontinued and thus reduce both patient burden and healthcare costs. Future periods of phenotyping may be required but will always be short term. This significantly increases the practical utility of HYPO-CHEAT over algorithms that require continuous CGM.

Fine tuning of risk profile and provision of targets. Hypos do not all pose the same risk: in those in which the glucose




Figure 2. Per cent time hypoglycaemic (per hour) plotted and forecast for two patients by Facebook Prophet. These graphs demonstrate the daily and weekly patterns in per cent time hypoglycaemic for two separate patients. Black dots indicate measured time hypoglycaemic per hour (median 0, hence the black line at this level); the dark blue line indicates the forecast of per cent time hypoglycaemic and the light blue area represents uncertainty intervals. The dark blue forecast line shows a clear spike each day with a broader, repeating pattern every seven days (indicated by red lines and confirmed by 7-day gaps between the dates on the x-axis), confirming a weekly pattern of percentage time hypoglycaemic. What can also be appreciated from this basic analysis in both patients is the difference between Friday and Saturday (higher risk) and Sunday to Thursday (lower risk).

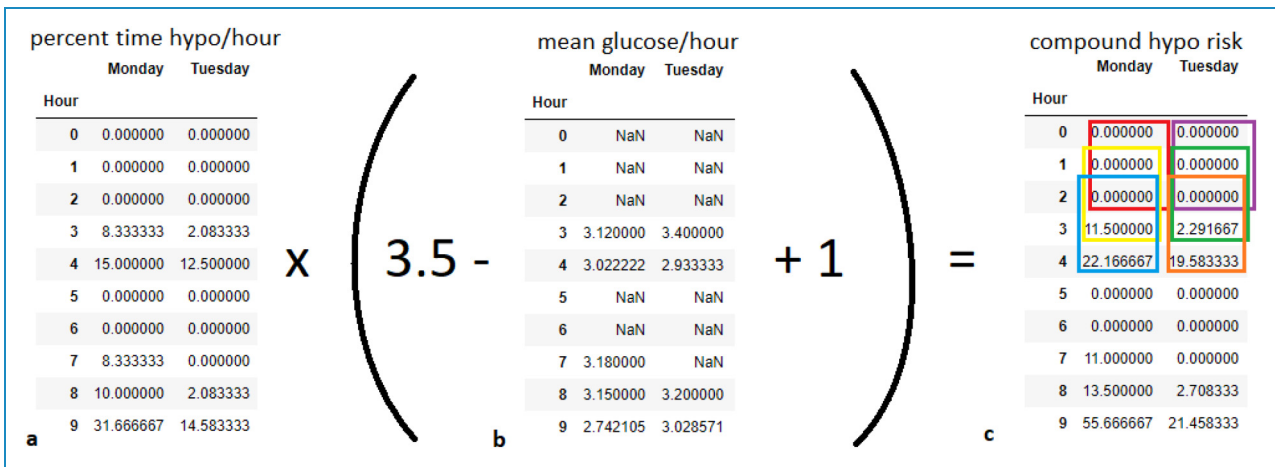


Figure 3. Visualisation of depth multiplier and the 3-hour buckets. (a) Shows the percentage time hypoglycaemic for each hour of each day. (b) Demonstrates the way that mean glucose per hour is used as a depth multiplier to result in a compound risk shown in (c) which shows coloured boxes which each represent a new bucket containing the composite hypoglycaemia risk for each containing hour. A mean is then computed for each bucket with the median hour providing the label for the bucket and its location on the heatmap. Thus, the periods of highest hypoglycaemia risk are ascertained, and an element of smoothing can be achieved across the heatmap.

drops to a lower level (deeper hypo), the risk to the patient is higher.⁴² Therefore, after discussions with the clinical team, it was felt to be appropriate to add a depth multiplier to the percentage time hypo for each hour of each day. After discussion, it was decided that $\{THRESHOLD - value + 1\}$ represented the most appropriate increase in severity for a depth multiplier. The way the multiplier works (Figure 3) is: for each hour containing any hypoglycaemia, the mean glucose value is calculated and subtracted from the threshold value. 1 is then added to prevent a situation of negative multipliers (e.g. $3.5 - 4.1 = -0.6$). This number is multiplied by the percentage time hypoglycaemic for that hour

(Figure 2). Thus, the degree to which the hypo drops below the hypoglycaemia threshold results in hypos with lower values being allocated more importance as the depth multiplier is a larger number.

HYPO-CHEAT subsequently evaluates every overlapping 3-hour bucket (e.g. 01:00–03:59H, 02:00–04:59H and 03:00–05:59H) within the 7-day window to ascertain the areas with the highest amount of hypoglycaemia (per cent time hypo \times depth multiplier) (Figure 3). Three-hour buckets were chosen after a trial of 1, 3 and 5-hour buckets in consultation with users. Odd numbers were used so that the central hour could be used as the bucket

label. One-hour buckets provided very little smoothing effect and patterns were difficult to detect as individual buckets did not reliably capture repeating behaviours. Five-hour buckets offered too much smoothing and it was harder to attribute hypos to single activities or behaviours. Three-hour buckets more reliably capture repetitive behaviours where times may vary slightly and be missed by 1-hour buckets (e.g. start time 07:50 1 week but 08:05 the next) but not be so large as to capture multiple behaviours. Thus, each bucket represents the risk from a limited activity profile and offers more utility on reflection of the causes of repeated hypoglycaemia.

A separate function ascertains which of the buckets contain a hypo every week and this is combined with the areas of highest risk to produce a set of buckets, which pose the greatest risk to patients based on a composite of time spent hypoglycaemic, depth of hypoglycaemia and frequency of recurrence. Each bucket is then compared with

corresponding buckets on other days (± 2 hours from the median). Areas of outstanding high risk compared with other days at the same time are added to the list of targets to be presented to the user. Areas of outstanding low risk are also highlighted if there is $>50\%$ less hypoglycaemia on 1 day than on all other days at the same time. This results in the significant outliers being reported to users who then can reflect on what is either causing or preventing hypos on that/those particular day(s) compared with other days at that time.

Any one-off hypos that represent a certain percentage of the highest repeat risk areas (set at 50% for our evaluation) are also reported. A cut-off of 50% was decided upon after evaluation of multiple levels from 20% to 80%. When the percentage was set very low there were too many one-off hypos to report and they distracted from the importance of the repeat hypos. At too high a level, no one-off hypos were ever reported. The level of 50% provided a point at




Figure 4. An example of a heatmap for Pt 5 generated over a period of 4 weeks (4w). This heatmap shows areas of repeat risk with darker areas representing higher compound risk (based on frequency and depth of hypos). Areas of repeat risk between 01:00 and 02:00 on Wednesday as well as between 08:00 and 09:00 on Thursday and Friday are quickly appreciated as the highest risk areas.

which the occasional hypo would be reported but very rarely more than two. A cut-off of two one-off hypos was also set for reporting to users. These are reported separately to repeat hypos and allow for reflection on a recent serious hypo but also to highlight this new area to try and avoid it becoming a repeat risk.

Visualisation of risk. A vital step of HYPO-CHEAT is the powerful way in which it presents the above information to users in the form of an easy-to-interpret heatmap. On a graph representing the week, this function allocates repeated periods of hypo risk with a red marker on the heatmap, where the shading indicates the compound severity of event as discussed above (e.g. Figure 4). The heatmap reads across in days from left to right and down in hours from midnight to 23:00. This order was chosen to mirror

popular online and digital calendars and thus aid quick interpretation by users.

Coloured markers are added to those areas in which hypos have occurred every (or most) week(s) since the previous analysis to draw even more attention to these areas of repeat risk and add to the ease of pattern identification for the user (Figure 5).

To avoid the loss of any information, any hypos which have not been seen to repeat are also represented on the same graph in blue (Figure 6). The maximum depth of colour for both red and blue is determined by the maximum compound hypo risk found in either category. This way, repeat areas (red) are likely to be primarily brought to the user's attention but any very severe one-off hypos (blue) will also be clear. This differentiation between repeat and one-off hypos allows the user to focus

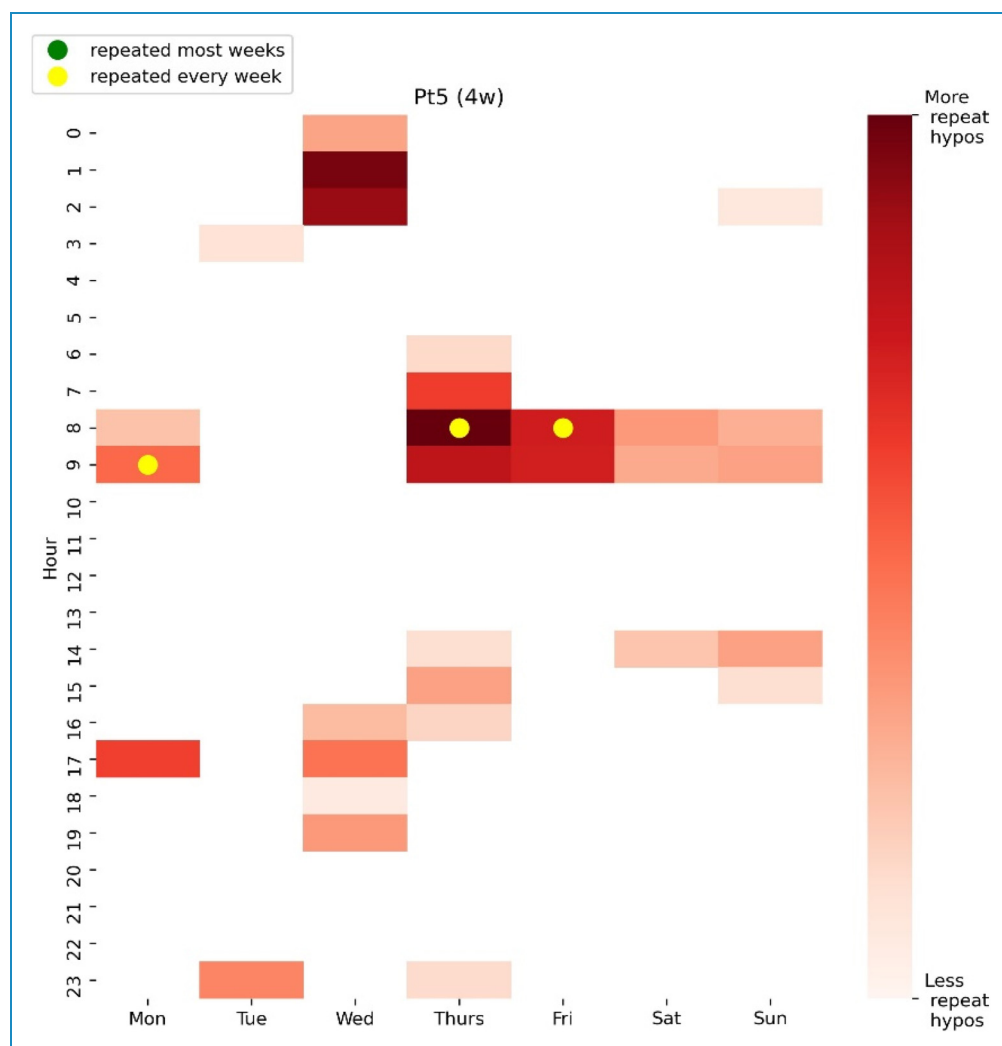


Figure 5. The addition of repeat markers to the heatmap for the same patient. Two of the three areas which previously appeared as the highest risk have received a repeat marker but that on Wednesday at 01:00–02:00 has not. This indicates that this risk area does not repeat every week and is based on fewer hypos. A new area has been identified as repeating every week (Monday 09:00). Thus, the areas upon which the user is asked to concentrate may have shifted away from Wednesday 01:00–02:00 and onto Monday in addition to Thursday and Friday around 08:00–09:00.

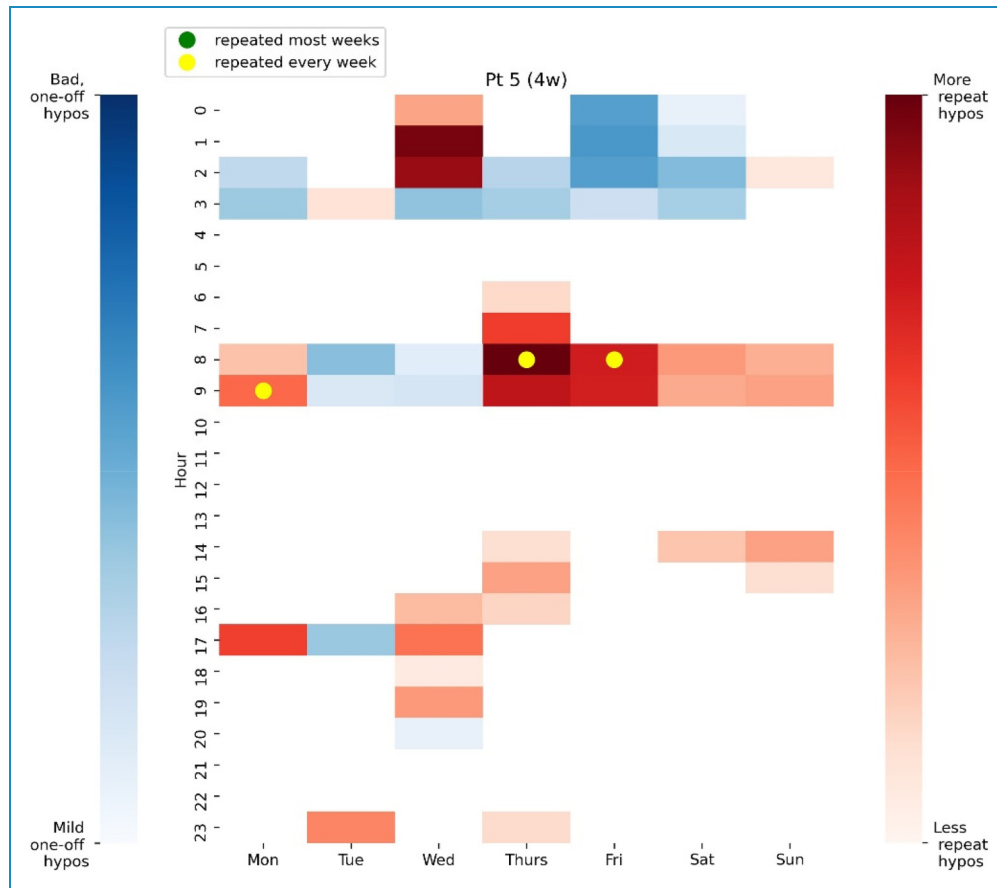


Figure 6. A fully complete heatmap for Pt 5. The addition of blue areas representing one-off hypos has contributed novel insight for the patient. Where before, one would have assumed that 01:00–03:00 was relatively safe other than on a Wednesday, it is now clear that the user often has hypos at this time but on variable days of the week. The addition of blue does not detract from the clear areas of highest risk on 08:00–09:00 on Monday, Thursday and Friday. However, the user can now understand that their patterns are primarily determined by time of day (08:00–09:00) but with an increase in this risk on certain days (Wednesday and Thursday) and a reduction on other days (Tuesday and Wednesday). This composite graph thus offers clear targets for focus but with additional information on which the user can reflect if they wish.

on areas of repeat (in red) without losing the important information about one-off hypos. In many cases, this can aid pattern finding as patients may have repeat hypos at a certain time on most days. Over a short monitoring period, a red risk area may not be created but a blue (and red) horizontal band of risk would be clear at that time of day across the week (e.g. the period 8:00–9:00 as shown in Figure 6). This would not be visible if one-off hypos were excluded, and a pattern could be missed.

The iterative component (assessing for improvement and providing feedback). HYPO-CHEAT is not a static algorithm but has iterative functionality to allow for the comparison of varying time periods. Once users have undergone an initial monitoring period with CGM to generate the first digital phenotype, then this will be adapted as they respond to targets and suggestions. Clear, text-based interpretation and feedback are also provided to facilitate understanding of what is indicated in the heatmap. HYPO-CHEAT is

designed to be used periodically to allow adequate time for significant changes in the weekly risk profile to have occurred. At each interaction with HYPO-CHEAT, the user is provided with an evaluation of how their risk profile has changed: Previous targets are assessed for the proximal outcome of fingerprick tests in these times and the more distal outcome of a change in composite hypoglycaemia risk. Users are praised for having undertaken fingerprick tests during target times and also if hypoglycaemia risk has decreased during targets. A comment is provided on whether the previous targets remain as targets or have been replaced by new time periods.

Severe, one-off hypos highlighted at the last interaction with HYPO-CHEAT are evaluated to assess if they have become repeats. If this is the case, then attention is drawn to this, and the user is asked to reflect on what has changed and why hypos in this period have worsened.

HYPO-CHEAT does have the capacity to function simply on fingerprick values after the initial period of

CGM phenotyping. In this situation, hypos will be cleared from the heatmap via targeted fingerprick tests that demonstrate repeat values above the threshold at times of the previous risk. New areas of hypoglycaemia will be added via a system of fingerprick checks at suggested times that will slowly roll forward. However, for this pilot study, to ensure as much information was obtained as possible, all assessments of change in hypos were based upon ongoing CGM data since the last HYPO-CHEAT interaction.

Evaluation

It was vitally important to us that HYPO-CHEAT was evaluated on real data and used with real patients in free-living conditions. Only by doing this could we truly evaluate if it was effective and identify areas for improvement and development. We thus designed a pilot project to test HYPO-CHEAT with 10 patients with CHI.

Methods

Patients with CHI were approached through the Northern Congenital Hyperinsulinism (NORCHI) highly specialised service at the Royal Manchester Children's Hospital (RMCH). Patients were eligible for inclusion if they:

- had a confirmed diagnosis of CHI;
- were receiving medication for the treatment of CHI and
- were under the age of 18 years.

Exclusion criteria were:

- already using a CGM device;
- likely to be able to stop treatment for CHI within the next 12 weeks and
- unable to attend 4 weekly appointments at RMCH.

Clinical leads were asked to suggest eligible patients for inclusion, and they were contacted via telephone by the research lead. Patient numbers were not based on any pre-determined power calculations but as per feasibility for a rare disease in a pilot study. As this system has not been tested previously, it was not known what reduction in hypoglycaemia could be expected so any power calculations would have been ill-informed and inaccurate.

All families attended the clinical research facility at RMCH where an initial CGM device (Dexcom G6) was attached and training on device usage was given. The study outline was explained to patients, as illustrated in Figure 7. Families were also provided with a Contour Next One glucometer and (in addition to targeted fingerpricks suggested later in the study by HYPO-CHEAT) were asked to undertake at least two fingerprick glucose tests per day to allow for an assessment of CGM device accuracy at the end of the study. For clinical safety, patients

were asked to perform a fingerprick test each time the CGM device reported hypoglycaemia when unblinded (see below).

For the first 4 weeks of the study, the CGM device was blinded to the users, so they were unable to see any readings. This period was designed as a baseline assessment of patients' glucose control and devices were blinded to ensure that patients were behaving and treating hypos as close to their normal as possible. In Week 4, the device was unblinded for the following 4 weeks, so patients had access to real-time readings, as well as the standard alarms on a Dexcom G6 (including low glucose and urgent low soon). In Week 8, the families were invited to a clinic appointment to use HYPO-CHEAT. They were provided with a single A4 PDF output of HYPO-CHEAT as well as being given a login to Dexcom Clarity software and time to review their data. Dexcom Clarity software is a standard interface provided by Dexcom for users to explore their glucose history and make changes to management. Initially, no explanation or instruction was given on how to use either. The study team compared HYPO-CHEAT's analysis with that of Dexcom Clarity (which has a dedicated section for patterns it has identified) and Facebook Prophet time series analysis [<https://facebook.github.io/prophet/>]. The ability to spot patterns of hypoglycaemia was compared between all three. These patterns and outputs from HYPO-CHEAT and Dexcom Clarity were discussed with families.

As Facebook Prophet is a forecasting algorithm and not designed for patient use, this output was not discussed with families but instead interpreted by the research team. Facebook Prophet is primarily designed to forecast a continuous variable and thus an initial use of Facebook Prophet where glucose values were used as the input did not yield interpretable results. Because we were not interested in the continuous variable, and to offer a fair comparison with our algorithm, we converted the input for Facebook Prophet to an hourly percentage of time spent hypoglycaemic rather than glucose values. This allowed the forecast to be of time spent hypoglycaemic (much like that of HYPO-CHEAT) rather than a continuous glucose value. This comparison offered an assessment of whether an ML algorithm would be able to identify more or less hypoglycaemia patterns than HYPO-CHEAT.

Finally, the CGM device was blinded again, and patients underwent a final 4 weeks of monitoring without access to real-time data but with the knowledge of the previous 8 weeks' data from HYPO-CHEAT. At the end of the monitoring period, analysis was undertaken to compare the time spent hypoglycaemic between each period (blinded – unblinded – reblinded). All periods were conducted under free-living conditions as patients went about their daily lives. No restrictions or requirements were placed upon patients other than to continue using CGM devices and provide regular SMBG checks.

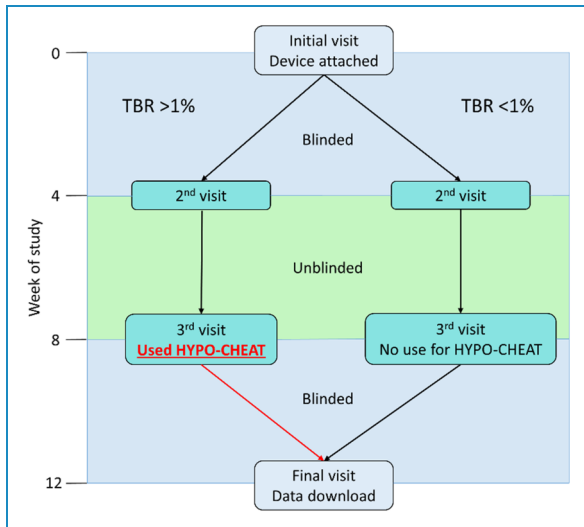


Figure 7. Outline of the study structure demonstrating the two groups separated by presence (left) or absence (right) of initial hypoglycaemia (TBR less than or more than 1%) in the first blinded period. Periods during which the CGM device was blinded (0–4 weeks and 8–12 weeks) are denoted in blue and the period of unblinded CGM use (4–8 weeks) is denoted in green. CGM: continuous glucose monitoring; TBR: time below range.

Results

The target for recruitment was 10 patients. Ultimately 12 families were approached. One family declined to take part due to the travel requirements of the study. The second patient declined, citing fear of painful CGM sensor insertion.

One patient withdrew from the study at the end of week eight citing problems with painful sensor insertion and overwhelming disturbance by CGM alarms. Of the remaining patients, there were five who demonstrated significant hypoglycaemia (TBR >1%) in the initial 4-week period (mean TBR 7.1%) and four who were almost completely free of hypoglycaemia (mean TBR 0.2%). HYPO-CHEAT was used with all nine patients but for those with no initial hypoglycaemia, there was almost no hypoglycaemia data to aggregate and analyse and thus the minimal utility of either HYPO-CHEAT or any other approach to prevent hypoglycaemia.

Pattern identification. The primary function of HYPO-CHEAT is a novel aggregation of CGM data into discrete weekly buckets and the presentation of this to patients in an actionable form. There was a clear pattern of hypoglycaemia




Figure 8. HYPO-CHEAT heatmap for Patient 2 (Pt 2) at the end of the unblinded period of 4 weeks (4w). Sunday around 1 pm stands out as a clear high-risk period. HYPO-CHEAT: **HY**poglycaemia-**P**revention-thrOugh-**CGM**-**HE**atmap-**T**echnology.




Figure 9. Facebook Prophet output for the same input as shown in Figure 8. This algorithm predicts a continuous variable which in this case is based on percentage time hypoglycaemic per hour. The y-axis reports a predicted percentage time hypoglycaemic. The obvious problem with a negative percentage time is clear but its use here is to simply spot the highest risk areas rather than absolute numbers. (a) Shows Sunday as the highest risk day and (b) shows 1 pm as the highest risk time. The risk seems to carry forward into Monday, which in reality is not true.

visible for all five patients with initial hypoglycaemia and these are discussed in turn. Full results are reported for the first patient (Patient 2) to orientate the reader to the method of result reporting. Individual results for other patients are provided in brief within this section and in full detail in Appendix 1.

Patient 2. Patient 2's family had not noticed any patterns in hypoglycaemia over the previous 4 weeks of blinded CGM use and 4 weeks of unblinded CGM use. HYPO-CHEAT highlighted a clear hotspot on Sunday around 1 pm for this patient (Figure 8). This hypo was repeated every week during this period. On being shown this output, the parents immediately identified that Patient 2 had a longer and later than usual nap on Sunday and, on reflection, realised that they did often wake up from this with a hypo.

When interrogating Facebook Prophet's analysis, Sunday around 1 pm is also identifiable as a period of high risk (Figure 9). However, from this output, it is not possible to ascertain how long this hypo normally lasts and the risk also appears to continue into Monday which, in reality, it does not (no repeats as shown in Figure 8 on Mondays). No patterns were identified by Dexcom Clarity (Figure 10).

Patient 3. Patient 3 and their family had not noticed any patterns in their hypoglycaemia. Following the review of

HYPO-CHEAT, clear behavioural causes for hypoglycaemia were identified. Dexcom Clarity reported no patterns and Facebook Prophet's output was non-specific. More details, as well as visualisations, are provided in Appendix 1.

Patient 5. The family of Patient 5 had spontaneously noticed early morning hypoglycaemia but no distinctive patterns. On review of HYPO-CHEAT, they identified later parental waking as contributing to worse early morning hypoglycaemia for Patient 5 on certain days. Dexcom Clarity and Facebook Prophet output, as well as more HYPO-CHEAT data and visualisations, are provided in Appendix 1.

Patient 8. Given a large amount of hypoglycaemia experienced by Patient 8, it is not surprising that they had not noticed any distinct patterns. After the review of HYPO-CHEAT, this family identified behavioural causes for the hotspots on two separate occasions. Dexcom Clarity and Facebook Prophet output, as well as more HYPO-CHEAT data and visualisations, are provided in Appendix 1.

Patient 9. Prior to review of HYPO-CHEAT, Patient 9 was adamant that no patterns of hypoglycaemia were present. However, after being shown their heatmap, Patient 9 identified clear behavioural causes for recurrent

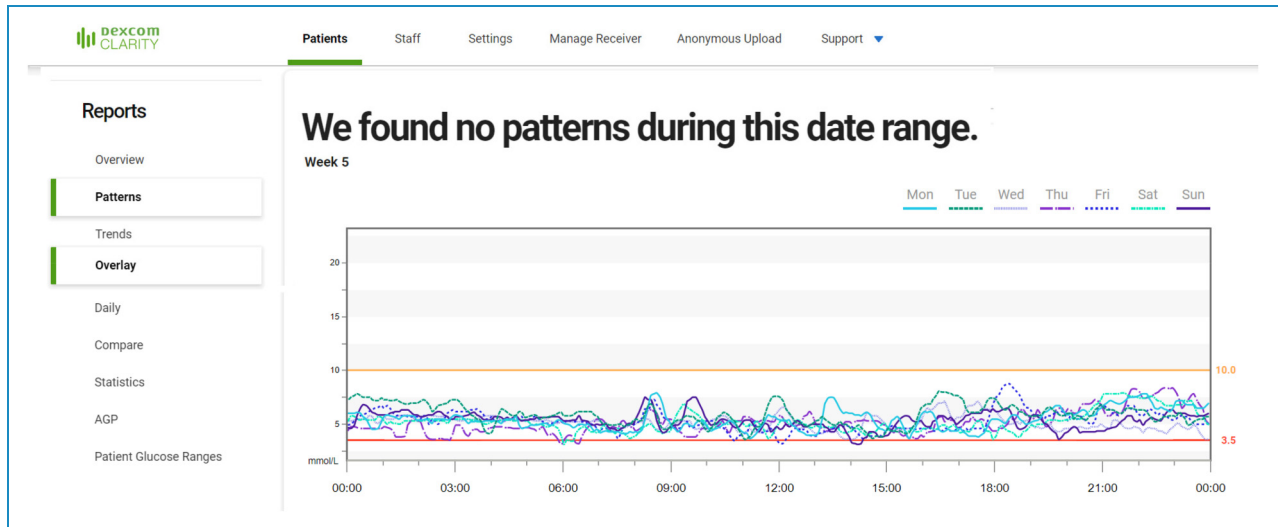


Figure 10. A composite Dexcom Clarity output for Patient 2 for the same 4-week period as shown in Figures 7 and 8. This screenshot (combined from two separate screenshots to display more information) displays the data available in the patterns tab (text indicating no patterns) and an example of data from the overlay tab (1 week shown at a time). Because the overlay only displays 1 week at a time, it cannot be appreciated that this patient suffers from more hypos on a Sunday afternoon despite the absolute consistency of this throughout the monitoring period (Figure 7).

hypoglycaemia. These, along with Dexcom Clarity and Facebook Prophet outputs, are detailed in Appendix 1.

Other patients. For the remaining patients, there was little to compare as the amount of time spent hypoglycaemic was so low that no patterns were identifiable (or likely present). Initial TBR was 0.2% for these patients and, while this did rise to 1.6% during the unblinded period, this was artificially elevated by a single patient who was unwell for a week, which resulted in a large number of hypos. Excluding this episode resulted in a mean TBR of 0.4% and thus the lack of repeatable patterns was maintained. Therefore, the analysis is restricted to those patients with initial hypoglycaemia and thus patterns that could be identified.

Change in hypoglycaemia following HYPO-CHEAT. Analysis of time spent hypoglycaemic (TBR) was the most distal outcome measured in this study and is summarised in Figure 7 and reported below. Proximal outcomes of behaviour change are reported in a separate paper focused on the behavioural change aspects of HYPO-CHEAT.⁴³

Of the nine patients who completed the study, five had initial hypoglycaemia (TBR >1%), with mean (range) TBR in the blinded (baseline) period of 7.1% (1.9–13.6). Following the use of HYPO-CHEAT, devices were reblinded for the final period and patients had no access to real-time data. During this period, TBR was 25% below baseline levels at 5.4% (2.8–11.0) and was significantly elevated by one patient who had spent a week in the hospital very unwell with recurrent hypoglycaemia.

For the four patients with almost no hypoglycaemia during the initial blinded (baseline) period, the mean TBR was 0.2%. While these patients were given access to HYPO-CHEAT, it was of little value as it is designed to prevent hypoglycaemia and for these patients, there was little to prevent. Throughout the course of the study; however, the TBR for these patients increased and during the final period was 3.2%. This increase in time hypoglycaemic is possibly reflective of patients having very tight control initially due to a feeling of being monitored and then possibly becoming complacent as the study progressed. It provides a useful comparison group to see how TBR changed when HYPO-CHEAT was not used and contrasts very starkly with the decrease in TBR seen in those who used HYPO-CHEAT.

User feedback on HYPO-CHEAT and Dexcom Clarity. Other than one parent, who found the volume of data on Dexcom Clarity to be helpful, all families reported that Dexcom Clarity provided too much information and it was very difficult to understand what their patterns were and what actions they should take to reduce their hypos. All families found utility in HYPO-CHEAT, preferring it to Dexcom Clarity as it was able to summarise their data and inform them if things were improving or not. Most importantly, families appreciated that following a quick read of HYPO-CHEAT's output they were able to reflect on behaviour and determine actions that could be taken to reduce the number and severity of hypos. The concentration of all the CGM info and its trends into a single A4 page was appreciated by families who liked the ability to quickly

understand how things were going and what still needed to be done.

Families particularly valued the separation of one-off and repeated hypos in the heatmap as this allowed them to consider behaviours that were contributing to repeat hypos. They valued the repeat green and yellow markers for the same reason. Finally, many of the families appreciated the provision of text-based interpretation and recommendations for improvement as it reduced the requirement for personal interpretation of a graph.

Discussion

We have outlined the development of HYPO-CHEAT, a novel approach to the detection, interpretation and visualisation of hypoglycaemia for patients with CHI which avoids the pitfalls faced by ML algorithms trying to predict a continuous glucose variable. At the base of the difference between HYPO-CHEAT and other algorithms is the use of the week as a discrete repeating unit. When days and hours are used as the repeating units, larger patterns are missed due to averaging effects. In most people, each day will not be the same as the others in the week and thus a daily repeating pattern is of limited use in most cases. This can be seen in the results above for HYPO-CHEAT, where only one out of five patients had a daily repeating pattern. All other patients had a weekly repeating pattern where hypos were clearly linked to behaviours performed on certain days of the week.

We compared HYPO-CHEAT's ability to evaluate CGM data and identify hypoglycaemia patterns with that of Dexcom Clarity software and Facebook Prophet time series forecasting algorithm. In all the patients who demonstrated significant TBR, HYPO-CHEAT was able to identify patterns of hypoglycaemia risk by interpreting each day of the week as a separate entity. In every case, the identified patterns correlated with the lived patient experiences and, thanks to HYPO-CHEAT's simple visualisations and text-based explanations, families were able to identify specific behaviours causing, or contributing towards these areas of risk. More importantly, this enabled them to take action to limit further events, potentially leading to improved outcomes in the long term for patients. In each of these cases, the families had not identified these patterns without the visualisation and prompting provided by HYPO-CHEAT.

Dexcom Clarity software was unable to identify any patterns of hypoglycaemia. Dexcom Clarity's inability to identify these patterns lies in its assumption that all days are the same and thus, if hypos are happening at different times on different days (as is to be expected when hypos are linked to behaviours such as eating, fasting and exercising), these are not identified as a pattern as they do not repeat frequently enough.

Facebook Prophet time series forecasting algorithm offered variable results in the successful identification of hypoglycaemia risk patterns. When a hypo was only seen at one time of day (such as Patients 2 and 5), it was possible to ascertain this time of risk, as well as the days on which this was worst, from the algorithm's output. However, if differing days showed hypos at different times (such as in Patients 3, 8 and 9) then the averaging effect of this algorithm meant that interpretation of the output either resulted in an incomplete picture (Patient 9 showed no Sunday afternoon risk) or a completely incorrect picture (Patient 3's output appeared to show Monday as a high risk at three times of day despite no hypos recorded on a Monday). Facebook Prophet time series forecasting was used to offer a comparison to HYPO-CHEAT and assess the differing abilities of the algorithms to identify patterns. However, Facebook Prophet is not designed for patient use and its output is therefore often difficult to interpret for non-expert users. As such, HYPO-CHEAT offers a clear advantage to patients.

A further advantage of HYPO-CHEAT is that, as hypos are visualised over the week, they can be associated with specific behaviours and those behaviours can be adapted. All patients were able to identify the behaviours contributing to their hypoglycaemia hotspots and thus had targets to action. This is cemented by HYPO-CHEAT providing each patient with text-based explanations of three targets for improvement before the next review and thus relieving the patient of the need to interpret graphs and charts if they do not want/are not able to do so. Neither Dexcom Clarity nor Facebook Prophet provides the relevant information that patients require to make decisions around their hypoglycaemias.

A proximal outcome of this pattern recognition was the improvement in TBR in those five patients who showed initial significant TBR and thus used HYPO-CHEAT for its intended purpose. The fact that TBR increased in those who did not use HYPO-CHEAT (due to very minimal initial hypoglycaemia) strengthens the argument that it is the use of HYPO-CHEAT which allowed patients to take control of their glucose profiles and prevent hypos from occurring. This is the very first demonstration of a CGM-informed algorithm effectively reducing hypoglycaemia in patients with CHI. To the best of our knowledge, this is also the only hypoglycaemia prevention algorithm not reliant on continuous CGM that has been shown to effectively reduce hypoglycaemia in free-living conditions.

HYPO-CHEAT has also been designed with cost saving in mind. CGM is expensive and not available to the majority of patients worldwide.¹⁴ For this pilot, HYPO-CHEAT adapted each patient's heatmap based on CGM, but the final outcomes were improved despite no access to CGM. Future steps will extend HYPO-CHEAT further to include prompts and remove the requirement for long-term CGM to remain iterative and up to date. Once HYPO-CHEAT

has built a stable phenotype of a patient's hypoglycaemia profile, CGM will be removed, and patients will be directed to undertake fingerprick glucose testing at particular times of day corresponding to periods of high risk. This will interact with and update the heatmap and feedback and targets will be provided. Intermittently, patients can undergo a period of CGM-based re-phenotyping to ensure the model is as up to date as possible. This method will generate huge cost savings as well as being available to a wider range of communities and healthcare settings rather than being restricted to economically advantaged healthcare systems. In the future, we intend to expand this approach to work with patients with T1DM and overcome the limitations currently in place to prevent this use (see the 'Limitations' section).

Limitations

HYPO-CHEAT addresses a lot of the issues inherent in preventing hypoglycaemia using ML and long-term CGM. However, there are inevitable limitations to our approach. The first is that, because HYPO-CHEAT is built around patterns, this model will not alert patients to impending and unexpected hypoglycaemia outside of recurrent behavioural patterns. This is a limitation that, it could be argued, is not faced by acute, continuous variable prediction algorithms. However, we believe that, with sufficient engagement with HYPO-CHEAT, patients will learn about behavioural triggers of hypoglycaemia and thus, unexpected hypos will become less common as patients change their behaviour appropriately. Ultimately, it is assumed that HYPO-CHEAT will offer a complementary approach to that of recursively updated dynamic prediction models (where ongoing CGM is available) so that patients will benefit from both behavioural pattern information as well as acute predictions. In situations where access to CGM is limited, HYPO-CHEAT offers a practical approach to maximise utility across a large patient group.

The second limitation of HYPO-CHEAT is that it still relies on a period of relatively expensive CGM wear to establish an initial digital phenotype. While the requirement for CGM is much lower than most alternative models, it is not zero and will potentially prevent the use of this approach in very low-income settings. Alternative models all require real-time CGM (live reporting of values) to function, but HYPO-CHEAT could be initially informed using non-real-time CGM (glucose measured continuously but only seen by the user when the device is scanned) such as the Freestyle Libre and thus make further cost savings.

Our evaluation of HYPO-CHEAT is somewhat limited by the lack of a true comparator group. While the patients with no initial hypoglycaemia did offer a useful comparison, they were not randomly selected, and they did have access to HYPO-CHEAT if desired (although no patterns would be recognised on such limited hypoglycaemia data). Future studies should randomise patients to

HYPO-CHEAT or alternatives such as simple CGM use; CGM use with clinician review and CGM use with patient review of Dexcom Clarity reports. Longer studies are also required to assess if TBR remains below baseline or if this effect reduces, or indeed increases, with time.

Finally, there is a limitation in the wider application of HYPO-CHEAT as it currently stands. HYPO-CHEAT is designed for patients with CHI who have recurrent hypoglycaemia but do not have to worry about high blood glucose (hyperglycaemia). If HYPO-CHEAT were to be extended to use in conditions such as diabetes, where hypo and hyperglycaemia are considerations, then significant adaptations would be required.

Conclusion

Current methods of hypoglycaemia prevention focus on the prediction of a continuous glucose value via ML and continuous CGM. Due to the multiple problems and reactive approach of this method, it is yet to show any real reduction in hypoglycaemia for patients in free-living conditions. HYPO-CHEAT is a novel approach to hypoglycaemia prevention that uses distinctive algorithmics to aggregate CGM data into discrete buckets and generate actionable visualisations and targets designed to proactively reduce hypoglycaemia for patients in the real world. A pilot of HYPO-CHEAT in 10 patients shows significantly improved performance over existing algorithms in the detection of patterns of hypoglycaemia that patients can act upon and thus have an effect. This pilot also shows the first reduction in hypoglycaemia using an algorithm not reliant on continuous CGM. Because our approach is both effective and cheap, HYPO-CHEAT has the potential to make a significant difference to patients in a short time frame. Related work describes the outcomes from the use of HYPO-CHEAT in terms of the more proximal outcomes of behaviour change.⁴³ Work is already underway to add further functionality for HYPO-CHEAT to be used with simple fingerprick monitoring along with more iterative analysis and assessment of and adaption to irregular changes in behaviour. Prompt and nudge technologies will be incorporated into future versions of HYPO-CHEAT and new patient groups will be engaged to assess the utility in varying pathologies. Finally, it is envisioned that, rather than an alternative approach, HYPO-CHEAT will be used as a complementary system alongside a recursive predictive algorithm and offer the best of both worlds.

Acknowledgements: The authors would like to thank Sumera Ahmed, Elaine O'Shea and Sarah Worthington for their assistance in the practical aspects of running the clinical trial of HYPO-CHEAT.

Declaration of Conflicting Interests: The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Contributorship: CW designed and built the analytics required for the project. CW designed and ran the clinical trial of HYPO-CHEAT. CW wrote the manuscript. All authors provided expert input into the design and running of the study, reviewed and edited the manuscript and approved the final version.

Ethical approval: This work received full ethical approval from the HRA (IRAS project ID 268245) and REC (reference 07/H1010/88).

Funding: The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the Health Innovation Manchester (grant number Momentum Grant).

Guarantor: CW.

ORCID iD: Chris Worth  <https://orcid.org/0000-0001-6609-2735>

References

1. Worth C, Yau D, Salomon Estebanez M, et al. Complexities in the medical management of hypoglycaemia due to congenital hyperinsulinism. Vol. 92, *Clinical endocrinology*. Wiley; 2020. p. 387–395.
2. Worth C, Tropeano Y, Gokul PR, et al. Insight into hypoglycemia frequency in congenital hyperinsulinism: evaluation of a large UK CGM dataset. *BMJ Open Diabetes Res Care* 2022; 10: e002849. Available from: <https://pubmed.ncbi.nlm.nih.gov/35675953/>.
3. Kanaka-Gantenbein C. Hypoglycemia in childhood: Long-term effects. In: *Pediatric Endocrinology Reviews*. 2004. pp. 530–536.
4. Worth C, Al Hashmi L, Yau D, et al. Longitudinal auxological recovery in a cohort of children with hyperinsulinaemic hypoglycaemia. *Orphanet J Rare Dis* 2020; 15: 1–9.
5. Lord K, Radcliffe J, Gallagher PR, et al. High risk of diabetes and neurobehavioral deficits in individuals with surgically treated hyperinsulinism. *J Clin Endocrinol Metab* 2015; 100: 4133–4139.
6. Lundkvist J, Berne C, Bolinder B, et al. The economic and quality of life impact of hypoglycemia. *Eur J Heal Econ* 2005; 6: 197–202. Available from: <https://pubmed.ncbi.nlm.nih.gov/15761775/>.
7. Eljamel S, Griffiths A, Evans J, et al. The burden of congenital hyperinsulinism in the United Kingdom: A cost of illness study. *Orphanet J Rare Dis* 2018; 13.
8. Thabit H, Prabhu JN, Mubita W, et al. Use of factory-calibrated real-time continuous glucose monitoring improves time in target and HbA1c in a multiethnic cohort of adolescents and young adults with type 1 diabetes: the MILLENNIALS study. *Diabetes Care* 2020; 43: 2537–2543.
9. Worth C, Dunne M, Ghosh A, et al. Continuous glucose monitoring for hypoglycaemia in children: perspectives in 2020. *Pediatr Diabetes* 2020 Aug; 21(5): 697–706.
10. Oviedo S, Vehí J, Calm R, et al. A review of personalized blood glucose prediction strategies for T1DM patients. *Int J Numer Method Biomed Eng* 2017; 33.
11. Worth C, Harper S, Salomon-Estebanez M, et al. Timing of Hypoglycaemia in Patients with Hyperinsulinism (HI): Extension of the Digital Phenotype. In: *Hormone Research in Paediatrics* [Internet]. 2021. p. 98. Available from: <https://www.karger.com/DOI/10.1159/000518849>.
12. Dave D, DeSalvo DJ, Haridas B, et al. Feature-Based Machine Learning Model for Real-Time Hypoglycemia Prediction. *J Diabetes Sci Technol* 2020.
13. Worth C, Harper S, Salomon-Estebanez M, et al. Clustering of hypoglycemia events in patients with hyperinsulinism: extension of the digital phenotype through retrospective data analysis. *J Med Internet Res* 2021; 23: 1–14.
14. Bila R, Varo R, Madrid L, et al. Continuous glucose monitoring in resource-constrained settings for hypoglycaemia detection: looking at the problem from the other side of the coin. *Biosensors* 2018; 8: 2–7.
15. Pérez-Gandía C, García-Sáez G, Subías D, et al. Decision support in diabetes care: the challenge of supporting patients in their daily living using a Mobile glucose predictor. *J Diabetes Sci Technol* 2018; 12: 243–250.
16. Mujahid O, Contreras I and Vehí J. Machine learning techniques for hypoglycemia prediction: trends and challenges. *Sensors (Switzerland)* 2021; 21: 1–21.
17. Turksoy K, Bayrak ES, Quinn L, et al. Hypoglycemia early alarm systems based on multivariable models. *Ind Eng Chem Res* 2013; 52: 12329–12336. Available from: <https://pubs.acs.org/doi/full/10.1021/ie3034015>.
18. Xie J and Wang Q. Benchmarking machine learning algorithms on blood glucose prediction for type 1 diabetes in comparison with classical time-series models. *IEEE Trans Biomed Eng* 2020; 67: 3101–3124.
19. Jensen MH, Dethlefsen C, Vestergaard P, et al. Prediction of nocturnal hypoglycemia from continuous glucose monitoring data in people with type 1 diabetes: a proof-of-concept study. *J Diabetes Sci Technol* 2019.
20. Vehí J, Contreras I, Oviedo S, et al. Prediction and prevention of hypoglycaemic events in type-1 diabetic patients using machine learning. *Health Informatics J* 2020 Mar; 26(1): 703–718. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/31195880>.
21. Turksoy K, Kilkus J, Hajizadeh I, et al. Hypoglycemia detection and carbohydrate suggestion in an artificial pancreas. *J Diabetes Sci Technol* 2016; 10: 1236. Available from: <https://pmc/articles/PMC5094335/>.
22. Kodama S, Fujihara K, Shiozaki H, et al. Ability of current machine learning algorithms to predict and detect hypoglycemia in patients with diabetes mellitus: meta-analysis. *JMIR Diabetes* 2021; 6.
23. Diouri O, Cigler M, Vettoretti M, et al. Hypoglycaemia detection and prediction techniques: a systematic review on the latest developments. *Diabetes Metab Res Rev* 2021 Oct; 37(7): e3449
24. Tyler NS and Jacobs PG. Artificial intelligence in decision support systems for type 1 diabetes. *Sensors (Switzerland)* 2020; 20: 1–26.

25. Lee S, Kim J, Park SW, et al. Toward a fully automated artificial pancreas system using a bioinspired reinforcement learning design: in silico validation. *IEEE J Biomed Heal Informatics* 2021; 25: 536–546.
26. Cryer PE. Minimizing hypoglycemia in diabetes. *Diabetes Care* 2015; 38: 1583–1591. Available from: <https://pubmed.ncbi.nlm.nih.gov/26207052/>.
27. Breton MD, Patek SD, Lv D, et al. Continuous glucose monitoring and insulin informed advisory system with automated titration and dosing of insulin reduces glucose variability in type 1 diabetes Mellitus. *Diabetes Technol Ther* 2018; 20: 531–540.
28. Tyler NS, Mosquera-Lopez CM, Wilson LM, et al. An artificial intelligence decision support system for the management of type 1 diabetes. *Nat Metab* 2020; 2: 612–619. Available from:
29. Skrøvseth SO, Årsand E, Godtliebsen F, et al. Data-Driven personalized feedback to patients with type 1 diabetes: a randomized trial. *Diabetes Technol Ther* 2015; 17: 482–489.
30. Liu C, Avari P, Leal Y, et al. A modular safety system for an insulin dose recommender: a feasibility study. *J Diabetes Sci Technol* 2020; 14: 87–96. Available from: <https://pubmed.ncbi.nlm.nih.gov/31117804/>.
31. Lind M, Ólafsdóttir AF, Hirsch IB, et al. Sustained intensive treatment and long-term effects on HbA1c reduction (SILVER Study) by CGM in people with type 1 diabetes treated with MDI. *Diabetes Care* 2021; 44: 141–149.
32. Davey RJ, Low C, Jones TW, et al. Contribution of an intrinsic lag of continuous glucose monitoring systems to differences in measured and actual glucose concentrations changing at variable rates in vitro. *J Diabetes Sci Technol* 2010; 4: 1393–1399.
33. Buckingham B, Block J, Burdick J, et al. Response to nocturnal alarms using a real-time glucose sensor. *Diabetes Technol Ther* 2005; 7: 440–447.
34. Perri A, Giordano L, Corsello M, et al. Continuous glucose monitoring (CGM) in very low birth weight newborns needing parenteral nutrition: validation and glycemic percentiles. *Ital J Pediatr* 2018; 44: 1–6.
35. Tomotaki S, Toyoshima K, Shimokaze T, et al. Reliability of real-time continuous glucose monitoring in infants. *Pediatr Int* 2019.
36. Alsaffar H, Turner L, Yung Z, et al. Continuous flash glucose monitoring in children with congenital hyperinsulinism; first report on accuracy and patient experience. *Int J Pediatr Endocrinol* 2018; 2018.
37. Rayannavar A, Elci OU, Mitteer L, et al. Continuous glucose monitoring systems: are they useful for evaluating glycemic control in children with hyperinsulinism? *Horm Res Paediatr* 2020; 92: 319–327.
38. Worth C, Dunne MJ, Salomon-Estebanez M, et al. The hypoglycaemia error grid a UK-wide Consensus on CGM accuracy assessment in hyperinsulinism. *Submiss to Front Endocrinol* 2022.
39. Danne T, Nimri R, Battelino T, et al. International consensus on use of continuous glucose monitoring. *Diabetes Care* 2017; 40: 1631–1640. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/29162583>.
40. McCarthy S. Weekly patterns, diet quality and energy balance. *Physiol Behav* 2014; 134: 55–59. Available from: <http://doi.org/10.1016/j.physbeh.2014.02.046>.
41. Nordman M, Matthiessen J, Biloft-Jensen A, et al. Weekly variation in diet and physical activity among 4-75-year-old Danes. *Public Health Nutr* 2020; 23: 1350–1361.
42. Rozance PJ and Hay WW. Hypoglycemia in newborn infants: features associated with adverse outcomes. *Biol Neonate* 2006; 90: 74–86.
43. Worth C, Nutter PW, Salomon-Estebanez M, et al. The behaviour change behind a successful pilot of hypoglycaemia reduction with HYPO-CHEAT. *Second Rev J Biomed Informatics* 2022.

Appendix

Appendix 1

Patient 3. Patient 3 and their family had not noticed any patterns in their hypoglycaemia during the blinded or unblinded periods of monitoring. HYPO-CHEAT showed a clear hotspot on Saturday morning at 10 am and a lesser hotspot on Sunday early evening (Figure 11). When reviewing this, the family realised that Patient 3 always woke later on a Saturday compared to every other day; Monday-Friday they had school and on Sunday they played football in the morning. Sunday morning football was identified as not only protective against morning hypoglycaemia (because they had to get up early) but likely contributory towards the afternoon hypoglycaemia.

It is impossible to appreciate the different time of day posing the highest risk on different days from Facebook Prophet analysis and Monday is attributed risk which in reality does not exist (Figure 12). No patterns were identified by Dexcom Clarity.

Patient 5. When asked if Patient 5 had any patterns to their hypoglycaemia, their family stated that they had some hypos in the morning but no pattern. HYPO-CHEAT showed that Patient 5 regularly had hypos every morning of the week apart from Sunday (Figure 13). This was clearly worse on Friday and Saturday mornings. On reflection, Patient 5's parents acknowledged that they normally woke later on Friday and Saturday and thus they received their breakfast later on those days and this was likely contributing to the high incidence of hypoglycaemia on those days. Sunday morning was identified by HYPO-CHEAT as an outlier and parents were asked to reflect on what they did differently on either Sunday morning or Saturday night.

Facebook Prophet's algorithm identified Saturday at around 7 am as a risk spot but attributed similar risk to Friday and Sunday (Figure 14) which is misleading and inaccurate. Dexcom Clarity software did identify a pattern but this was of hyperglycaemia (high blood glucose) rather than hypoglycaemia (Figure 15). This demonstrates one of the issues inherent in using software not designed for a single and specific purpose (e.g. to prevent hypoglycaemia). Dexcom Clarity did not identify any patterns of hypoglycaemia during this period.




Figure 11. HYPO-CHEAT heatmap for Patient 3 (Pt 3) at the end of the unblinded period. Saturday morning is the highest risk, closely followed by Sunday afternoon.

HYPO-CHEAT: **HY**poglycaemia-**P**revention-thr**O**ugh-**CGM**-**H**eatmap-**T**echnology.

Patient 8. Patient 8's family had not identified any patterns of hypoglycaemia over the previous unblinded monitoring period. Given the large number of hypos experienced by this patient (11–14% of the time) it was understandable that parents had not identified patterns. However, some patterns were visible when viewing HYPO-CHEAT's output (Figure 16). The standout areas of risk were Tuesday around 4pm and very early on Sunday morning. When shown this output, the patient themselves identified that they did physical education (PE) lessons at school on Tuesday afternoon and then walked home and this was likely causing this hotspot. Then patient's parents identified that mealtimes were very different on a Saturday, and this was likely causing the hypos overnight Saturday into Sunday.

The Facebook Prophet output is hard to interpret (Figure 17) and certainly does not show Tuesday at 4pm and Sunday at 1am as clear hotspots. No patterns were identified by Dexcom Clarity.

Patient 9. Patient 9 was very clear that they had not spotted any patterns of hypoglycaemia as they had simply been irritated by the low alert alarms and had thus not noticed the patterns of when they occurred. HYPO-CHEAT identified the main areas of risk for Patient 9 as overnight on Saturday and Tuesday with a separate and lesser spot on Sunday evening (Figure 18). After being shown this output, Patient 9 identified that overnight hypos were very likely associated with athletic training the night before as well as possible alcohol consumption. They also commented that the Sunday afternoon hypo was similarly following athletic activities.

Facebook Prophet correctly showed overnight as the higher risk period and Saturday as the worst day with a smaller peak on Tuesday. However, due to averaging effects, there was no identification of the Sunday afternoon risk period (Figure 19). No patterns were identified by Dexcom Clarity.




Figure 12. Facebook Prophet output for the same input as shown in Figure 11. To try and capture all areas of risk, the predictions are not helpful as the algorithm cannot appreciate the differing times of risk (a) between days (b) and thus has multiple time spikes and averages the highest risk to a day which in reality contains no hypos (Monday).




Figure 13. HYPO-CHEAT heatmap for Patient 5 (Pt 5) at the end of the unblinded period. Early morning is a clear period of high risk which is increased on Friday and Saturday and minimised significantly on Sunday.

HYPO-CHEAT:

HYpoglycaemia-**P**revention-thr**O**ugh-**C**GM-**H**eatmap-**T**echnology.




Figure 14. Facebook Prophet output for the same input as shown in Figure 13. Early morning risk is identified (a) and Saturday is appropriately high risk (b). However, the increased risk of Friday and decreased risk of Sunday is not clear.

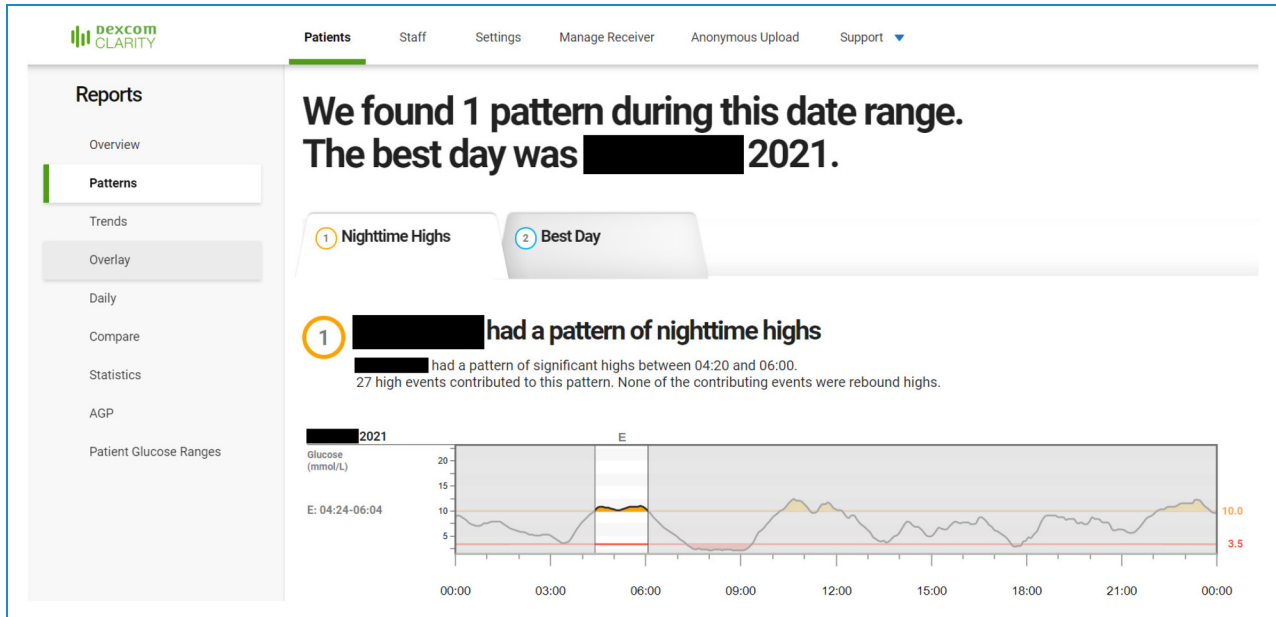


Figure 15. Dexcom Clarity pattern tab for Patient 5 for the same 4-week period as Figures 13 and 14. For this patient, Dexcom Clarity did identify a daily pattern of high blood glucose but failed to identify any patterns of hypoglycaemia (the hypoglycaemia following high blood glucose can be appreciated on the day highlighted in this figure). Patient name and date have been redacted to preserve confidentiality.

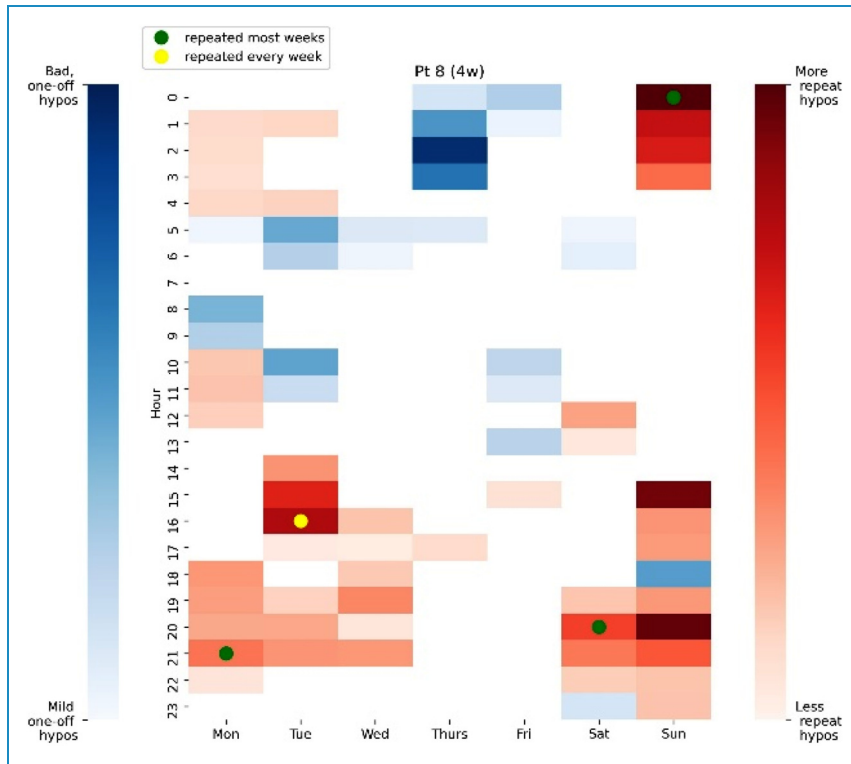


Figure 16. HYPO-CHEAT heatmap for Patient 8 (Pt 8) at the end of the unblinded period. While there is a lot of hypoglycaemia for this patient, HYPO-CHEAT identified Tuesday around 4 pm and Sunday around midnight as the highest risk areas and targets for focus. HYPO-CHEAT: **HY**poglycaemia-**P**revention-thr**O**ugh-**CGM**-**H**eatmap-**T**echnology.

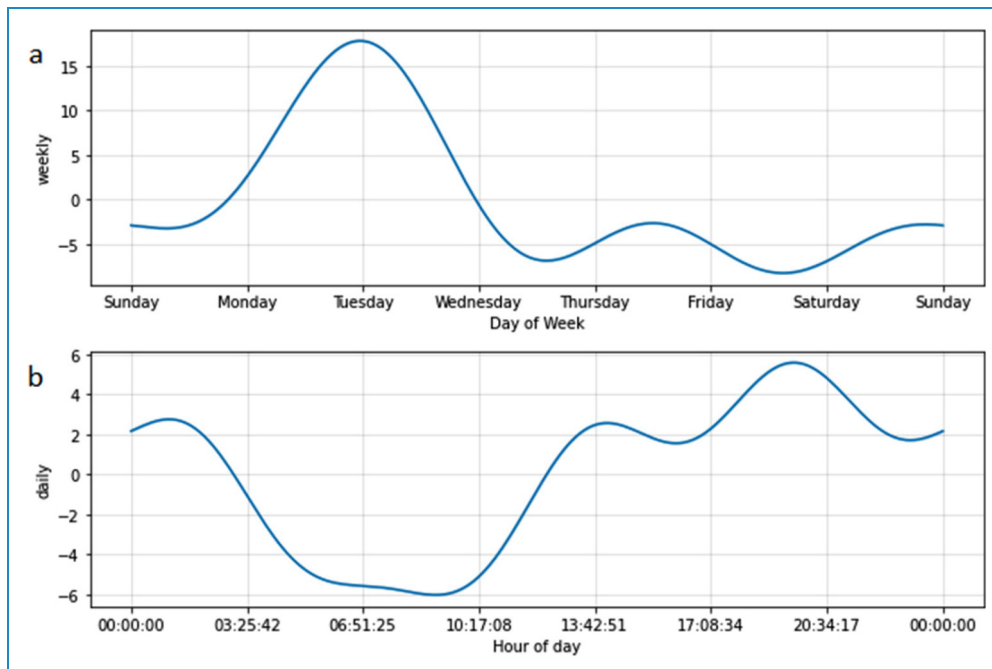


Figure 17. Facebook Prophet output for the same input as shown in Figure 16. Tuesday is identified as a high-risk day (a) but from the time profile (b) one would expect this risk to be highest around 8 pm (which, in reality, it is not). The high risk of Sunday is not identified at all.




Figure 18. HYPO-CHEAT heatmap for Patient 9 (Pt 9) at the end of the unblinded period. Saturday and Tuesday overnight were clear risk periods with a lesser period of risk on Sunday afternoon. HYPO-CHEAT: **H**Ypoglycaemia-**P**revention-thr**O**ugh-**C**GM-**H**Eatmap-**T**echnology.




Figure 19. Facebook Prophet output for the same input as shown in Figure 18. Overnight (a) and Sunday (b) are clearly shown as risk areas with a smaller peak on Tuesday. The Sunday afternoon risk is not identified at all.