
Proceedings Paper:
Hama, L., Ruddle, R.A., Abuzour, A.S. et al. (Accepted: 2025) Patient History Visualization for Structured Medication Reviews: A Design Study. In: To be confirmed. 19th International Joint Conference on Biomedical Engineering Systems and Technologies, 02-04 Mar 2026, Marbella, Spain. (In Press)

Patient History Visualization for Structured Medication Reviews: A Design Study

L. Hama[8,10], R.A. Ruddle[8,10], A.S. Abuzour[1,2], M. Abaho[3], A. Aslam[2,8], D. Bollegala[3], A. Clegg[1,2], M. Gabbay[3], F.S. Mair[5], M. O'Connell[6], S. Relton[2,8], E. Shantsila[3], M. Sperrin[6], S.A. Wilson[3], A.A. Woodall[3,4], I. Buchan[3], L.E. Walker[3]

[1]*Bradford Teaching Hospitals NHS Foundation Trust, Bradford, UK*

[2]*School of Medicine, University of Leeds, Leeds, UK*

[3]*University of Liverpool, Liverpool, UK*

[4]*Directorate of Mental Health and Learning Disabilities, Powys Teaching Health Board, Bronllys, UK*

[5]*School of Health and Wellbeing, University of Glasgow, Glasgow, UK*

[6]*University of Manchester, Manchester, UK*

[8]*Leeds Institute for Data Analytics, University of Leeds, Leeds, UK*

[10]*School of Computing, University of Leeds, Leeds, UK*

email:layik.hama@gmail.com

Keywords: Patient History Visualization, Deprescription Visualization.

Abstract: General practitioners and pharmacists conduct Structured Medication Reviews (SMRs) to optimise prescribing for people with multiple long-term conditions (MLTC), but electronic health record systems often present information in fragmented lists and tabs. We set out to design and validate chart-based visual summaries of patient history data that can support integrated dashboards for SMRs. Using a design-study methodology, we reviewed existing approaches to visualising electronic health records, conducted four mock SMRs, and derived a data abstraction linking clinicians' questions to patient attributes. Using visual encoding principles, we used 14 candidate chart types and sketched low-fidelity designs. A questionnaire was then used to ask eight clinicians and eleven visualization-literate researchers to rate how effectively each chart communicated its data. Across six combinations of data types, 11 chart types were consistently judged suitable for communicating key information. Finally, we implemented the outcomes of the study in a Python package using 2.1M records from the Clinical Practice Research Datalink (CPRD).

1 Introduction

General practitioners (GPs) and pharmacists conduct Structured Medication Reviews (SMRs) (Madden et al., 2022) as "a comprehensive and clinical review of a patient's medicines and detailed aspects of their health" (NHS England, 2021). SMRs involve consultation with patients, considering their medical history (diagnoses, symptoms, current medication, investigations, etc.) (Madden et al., 2022). In the UK medical history is accessed via an electronic health record (EHR) system such as EMIS (EMIS Health, 2023). Such systems provide detailed information about a patient's medical history but information that is needed to conduct SMRs in complex patients is not presented in a manner that is time efficient.

A wide variety of techniques have been applied to visualise EHRs of individual patients and cohorts of patients in previous research (Wang and Laramee, 2022; Dowding et al., 2015), but not for the purpose of conducting medication reviews. The present work follows a design study methodology (Sedlmair et al., 2012) to select candidate charts for visual histories of patient records for SMRs. Then, four mock patient SMRs were carried out to establish user requirements in terms of: (a) questions that GPs and or pharmacists typically want to answer in an SMR, and (b) the attributes of medical histories that are required. A design space was determined by mapping the attributes to the data types that different types of visualization support. Low-fidelity (paper and pencil) prototypes were used to evaluate the design choices and produce a shortlist of visualizations for groups of SMR questions.

This paper makes three contributions based on the methodology (Sedlmair et al., 2012) followed. First, problem characterisation, that is, analysing and data abstraction using mock SMRs. Second, it presents candidate visualizations for six combinations of data types. Third, a Python package that implements the outcomes (DynAIRx Project, 2025).

The rest of this work is structured as follows: It starts with an extensive background review of techniques (charts) used to visualise EHRs. Then it summarises the work done to understand the nature of SMRs using the four mock SMRs. Section 4 then describes the design space exploration approach and the outcomes and one round of validation of our design choices. The last section outlines the implementation of the design study.

2 Related Work

The present work aims to develop visual summaries of patient medical histories in the context of SMR. Expert input from clinicians who are part of the project team and focus groups with other clinicians shows that these summaries need to contain longitudinal information on the conditions, medications, and investigations of a patient.

Visual summaries also need to present the information in a much easier to digest form than the multiple tabs and click-heavy user interfaces of general-purpose EHR systems (Abuzour et al., 2023). To our knowledge, there has been no previous research on how to visualise patient histories specifically for SMRs. However, a considerable body of work has investigated dashboards for EHRs (Dowding et al., 2015) and EHR visualization techniques (Wang and Laramee, 2022). We approached the related work using a semi-systematic method (Snyder, 2019), combining two review papers with forward citation tracking. The first was a review by Dowding et al. (Dowding et al., 2015) and the second was a state-of-the-art review by Wang and Laramee (Wang and Laramee, 2022). The criteria used to select the publications from these sources were: the work uses patient-level or cohort EHR, and is developed for clinical use.

2.1 EHR Dashboards

The first observation from the publications in our literature review is that the technique appearing the most is the use of tables along with bar and line charts. Some dashboards combine tables with other visualization techniques such as bars, pie, and traffic light colours (Koopman et al., 2011; Stinson et al., 2012; Daley et al., 2013).

The second observation is that we can see some uncommon visualization techniques. For example, the use of what is called a "signal element" (Schulz et al., 2020) is designed to create dashboards using the Arden syntax (Hripcsak, 1994) which is a markup language for sharing medical language. An important concept in EHR data is a threshold or a range of thresholds for particular measurements. This explains the appearance of a "bar graph (bar meter)" (Kheterpal et al., 2018).

The third observation is the variety of health record measurements appearing in the various dashboards. Unsurprisingly, there are time (date/interval, range of days or duration), state, trend, events, and in many cases combinations of such measurements. Visualization of status could be of a test being tracked (due, done, etc.) (Stinson et al., 2012), or the status of an organ (Kheterpal et al., 2018; Calzoni et al., 2020).

The fourth observation is about rare techniques such as glyphs for patient visits (as noted by (Wang and Laramee, 2022)). One of the dashboards (Kheterpal et al., 2018) uses an outline of the human upper body organs as the main view of their dashboard with bar graphs of measurements such as temperature and various blood test measurements. A single icon traffic light combined with text values for various measurements is also used to indicate status (Kraus et al., 2018). Another visualization technique in the context of the Intensive Care Unit (ICU) monitoring dashboard is the use of icons to represent the trend of a particular measurement over time (Calzoni et al., 2020). The technique used is a single-value scatter plot that represents the measurement in both its state (normal/abnormal) and its quantity but also where the measurement lies within the ranges of clinical guidelines over time.

3 Understanding SMRs

Following the recommendations of the design study methodology (Sedlmair et al., 2012), the first task was to acquire a thorough understanding of SMRs guided by practicing clinicians. The goal was to understand the questions commonly asked during SMRs. As real consultations contain confidential patient data that were not available to us, four mock SMRs were performed to replicate the medication review process and capture representative questions.

3.1 Method

A clinician co-author secured permission to run an empty EMIS (Egton Medical Information System) (EMIS Health, 2023) instance (version 9.18.11) on a project-dedicated Windows laptop. EMIS is a widely used GP EHR platform, holding an estimated 56% share of UK GP clinics in 2018 (Kontopantelis et al., 2018). Due to license restrictions, screenshots of the EMIS interface, whether or not they contained mock data, could not be reproduced here. Figure 1 shows a typical layout of an EMIS view used in the mock SMRs simulating the medications tab.

Sequence	Drug / Dosage / Quantity	Usage Current / Average	Quantity	Repeat	Altered	Issue Dat	Issuer
A	Ensure Compact liquid (Flavour Not Specified)	3 daily, 21 x 125 ml		no			Walker, La 0 of 3
B	Ensure liquid (Flavour Not Specified)	2 x daily		no			Walker, La 0 of 3
C	Olanzapine 10mg tablets	One To Be Taken Each Day		yes	altered	45116	Walker, La 0 of 3
D	Atenolol 50mg tablets	One To Be Taken Each Day		yes		45118	Walker, Lauren (Dr)
E	Carbocisteine 375mg capsules	Two To Be Taken Three Times A Day		yes			Walker, Lauren (Dr)
F	Clopidogrel 75mg tablets	One To Be Taken Each Day	28 tablet	yes		44570	Walker, Lauren (Dr)
G	Furosemide 40mg tablets	One To Be Taken Each Morning	28 tablet	yes			Walker, Lauren (Dr)
H	Gabapentin 300mg capsules	One To Be Taken Three Times A Day	90 capsule	yes			Woodall, Alan (Dr)
I	Lisinopril 5mg tablets	One To Be Taken Each Day	28 tablet	yes			Woodall, Alan (Dr)
J	Metformin 500ma modified-release tablets	One To Be Taken Each Day With Food	28 tablet	yes			Woodall, Alan (Dr)
K	Paracetamol 500mg tablets	TWO TO BE TAKEN FOUR TIMES DAILY	100 tablet	yes		44378	Woodall, Alan (Dr)
L	Salamol 100micrograms/dose inhaler CFC free (Teva UK Ltd	One Or Two Puffs To Be Inhaled Up To Four	2 x 200 dose	yes			Woodall, Alan (Dr)
M	Sertraline 50mg tablets	One To Be Taken Each Day	28 tablet	yes			Woodall, Alan (Dr)
N	Simvastatin 40mg tablets	ONE TO BE TAKEN DAILY	28 tablet	yes			Woodall, Alan (Dr)
O	Temazepam 20mg tablets	ONE TO BE TAKEN AT NIGHT	28 tablet	yes			Woodall, A 0 of 3
P	Trimbow 87micrograms/dose / 5micrograms/dose / 9microg	Two Puffs To Be Inhaled Twice A Day	2 x 120 dose	yes			Woodall, Alan (Dr)
Q	Warfarin 1mg tablets	Take As Instructed By Anticoagulant Clinic	28 tablet	yes			Woodall, A 0 of 3
R	Insulin glargine 100units/ml solution for injection 10ml vials	AS DIRE, 6 vial		no			Walker, La 0 of 3
S	Otosporin ear drops (Phoenix Labs Ltd)	TWICE A DAY	10 ml	yes		21-11-202	Walker, La 0 of 3

Figure 1: The diagram shows the medication view of the EMIS instance used during the mock SMRs. The look and feel of the EMIS version used was similar to Microsoft Word 2007.

The sessions were carried out over two days at the University of Liverpool and recorded using Microsoft Teams on the designated laptop that was also used to access the EMIS instance. We estimate that the total time required to be four days of combined effort of data population and conducting the four mock SMRs.

The procedure for the mock SMRs involved playing the roles of both patient and doctor by the two clinician co-authors, who simulated patient interactions as realistically as possible. For two of the patients, the clinician was a consultant clinical pharmacologist and the other two were performed by a GP. SMRs are typically made up of two separate parts: preparation and actual review with the patient, the same was true for the mock SMRs. For these mock SMRs, the clinicians did the preparation part as "think-aloud" sessions so that their thoughts and actions could be recorded. Typically, according to our two clinician co-authors, GPs would approach SMRs slightly differently from how pharmacists would do them, for instance. Another key element of the process was the flow of looking at various patient records to perform the SMR. For this, we noted that the GP and the consultant pharmacologist would primarily inspect four tabs on the EMIS application both during preparation and review with the patient present: **consultations**, **medication**, **problems** and **investigations**. We call these four tabs the "main tabs" for the present work from here on.

The process of conducting mock SMRs was initiated by having two of our clinician co-authors populate the dedicated EMIS instance with data for four mock patients. These patients were labeled as patients 5, 8, 10, and 11, and the clinicians ensured that the data for each patient was realistic and varied in complexity.

An example of a patient profile is Patient 5. This was a young 33-year-old female with early type 2 diabetes, asthma, the start of acne vulgaris, polycystic ovary syndrome, and self-harm; she smokes and drinks heavily and resists medications, poorly concordant with monitoring of her chronic disease condition, and was becoming overweight. The patient had nine drugs listed in the "repeat" section with a total of 12 medication rows appearing in the medication tab, including inhalers, tablets such as metformin for diabetes, and analgetic medications (painkillers) with potential for addition. The GP reviewed the four main tabs in the EMIS application in the preparation part of the review. They then moved to review the medications tab. The GP then looked at the investigation tab and said they wanted "to see how things are controlled". One of the line charts the GP inspected there was the weight of the patient over time. The chart showed a weight range of 50KG to 105KG between 2004 and 2022 for the mock patient. The GP then looked at the records on the consultation tab, noting various records. One of these records was that the patient had been excluded from a certain target by the GP practice due to their lack of attendance and he wanted to discuss this with the patient.

Data for conditions, medications, investigations, and consultations for the four patients are summarised in Table 1. Not all rows under each of the consultations, problems, medications, and investigations are unique, they can be repetitive, and some entries may not be directly related to where the entry should be. The average number of rows in problems (conditions) and medications tabs were 23 and 13 respectively.

3.2 Analysis of the SMRs

In this section, the SMR content is analysed to understand and characterise the domain problem. Some of the statistics such as descriptive statistics related to the content, data, and durations of SMRs conducted are reported. The analysis bridges the gap from the previous background section by generating the data abstraction required to conduct the design

Table 1: The table summarises time spent and row counts of each of the four tabs on the EMIS instance for the four mock SMRs patients.

Patient ID	EMIS Tab	Time (secs)	Row count
5	Consultations	12	100+
	Problems	52	20
	Medication	144	12
	Investigations	311	100+
8	Consultations	0	0
	Problems	108	50
	Medication	832	20
	Investigations	136	100+
10	Consultations	0	0
	Problems	201	6
	Medication	181	5
	Investigations	152	100+
11	Consultations	100	100+
	Problems	35	16
	Medication	913	16
	Investigations	177	100+

study.

3.3 Content of the SMRs

Of the four mock SMRs, a total of 117 questions were extracted from the transcripts generated by Microsoft Teams recordings. The questions were first extracted using OpenAI ChatGPT4 (OpenAI, 2023) to analyse the transcripts using this prompt "these are mock patient-doctor consultations, extract all the questions asked by the doctors". This was followed by manually checking each of the 117 questions within the video recordings to avoid transcription and extraction errors. The questions were then grouped into these categories: conditions, medications, investigations, emotional/well-being, lifestyle, personal, and personal histories.

These categories were later presented to two clinician co-authors for expert review. They made some changes in the categorisation with the final categories being *active problems*, medications, investigations, emotional/well-being, lifestyle, *social history* and *medical history*. For the remainder of this work, we will refer to active problems as "conditions". This was followed by a review of question categorisation (which question belongs to which category on a spreadsheet) by the clinicians. The final categorisation of the questions showed that the highest number is 49 in the category of medications. The medical history category has the least number of questions. The present work will focus on the following three categories of questions: **conditions**, **medications** and **investigations**.

3.4 Data abstraction

Having looked at the content and duration of the SMRs, the questions were further analysed to determine the data required for each category: **conditions**, **medications** and **investigations**. The first step was to list what pieces of information are required by clinicians to answer each question within those categories. From this list, for each of the questions, the combination of information such as name only, name and date, or a combination of name, date and other details was extracted. We call this information data attributes. For instance, for the question "Have you been on the two for a long time?" the required data was recorded as "medications" and "dates".

The next step was to identify the data types of the information required by each question. The data types used are nominal (e.g., the names of conditions or medications), quantitative, temporal, and ordinal. These are based on a taxonomy (Elshehaly et al., 2018) which is based on the Vega (Satyanarayan et al., 2015) data types.

The final step was to group the questions into possible combinations of data attributes for each category of question. The final output of this data abstraction process is summarised in Table 2 and the full list of questions is at `https://github.com/DynAIRx/HI2026`.

3.5 Summary

Four mock SMRs were performed to understand and characterise the domain problem (SMRs) believed to be the first of its kind for the purpose of a visualization design study focused on patient history visualization. The transcipts of these mock SMRs were analysed and a list of questions was extracted. From these questions a data abstraction was generated that scoped the design study, which will be discussed next.

4 Visualization design

In this section, the approach and results of generating design choices for the domain problem (clinicians conducting SMRs) are described. Generating these choices involved a two-step process: exploring the design space for design choices and evaluating them.

Table 2: The table shows breaking down questions asked during mock SMRs to generic data attributes required for each category of questions. The present work considers "dose" to be a quantitative value.

Question	**Data Attributes**	**Question Category**
What about taking the medicines because you're on quite a few. How do you feel about those? Has anybody ever mentioned something called COPD to you? Did he end up getting a scan in the end or not?	Name	Medications Conditions Investigations
Have you been on the two for a long time? Have you had a recent blood test looking for hepatitis? Did that [diabetes] start later on?	Name, Date	Medications Investigations Conditions
I see that you're being put on some fairly strong painkillers. You don't feel very strong, no?	Name, Dose	Medications
Have you ever had any specialist pain advice?	Name, Date, Severity	Conditions
Is that[diabetes] controlled at the moment?	Name, Date, Quantity	Investigations
Here's your blood pressure?	Quantity	Investigations

4.1 Design space exploration

The exploration of the design space based on charts was carried out according to the design guidelines and principles within the visualization literature. In this step of the design study, how each of the chart design choices would be limited by guidelines such as those of Munzner (Munzner, 2014) and Spence (Spence, 2001) was checked. Specifically, visual encoding channels such as position, size, shape, etc. The context of magnitude (quantitative values) and identity (nominal values) of each chart option were checked and noted. For example, in the case of identity channels such as colour hue, a limit of seven bins is recommended (Munzner, 2014).

To generate chart design options for each of the combinations of data types, a comprehensive list of charts was prepared from the literature review as well as Data Visualization Survey's 2021 chart lists (DVS, 2021). For each chart option in the list, the possible channel was recorded to encode each of the data type combinations (see Section 3.4). For instance, the height of a bar chart could be used to encode quantitative as well as temporal data types. We then filtered the charts based on the matching data types of the data type combinations for the three categories of SMR data: conditions, medications and investigations.

The final list of chart options for the categories of mock SMR questions was 14 (see Figure 3. Each chart choice in this table appears in the two review publications (Dowding et al., 2015; Wang and Laramee, 2022)) or the original publications cited by them.

This step was followed by three rounds of pencil and paper sketches excluding the fourth round of finalising the design options for evaluation over several weeks. In each iteration, the focus was on a combination of data types such as "NQ" (nominal, quantitative), and which category of questions the data would be visualized under (e.g. conditions). Pencil and paper sketching is more flexible and provides unrestricted space (Seitamaa-Hakkarainen and Hakkarainen, 2000) for being creative without restrictions compared to writing code from the beginning. Figure 2 shows two sets of design choices for the two combinations of data types: nominal, quantitative (N, Q) and nominal, temporal and ordinal (N, T, O).

4.2 Design choice evaluation

4.2.1 Method

The participants for the evaluation were chosen to be both domain experts (co-authors) and members of the research project team. All of the team were invited to participate and 18 responded.

In terms of materials, the final designs were generated for the list of charts for each of the combinations. These were drawn within equal-sized areas on A4 sheets and scanned. Each chart was populated with three rows of data. The same quantitative values were repeated to create a list of 10 in the case of box, dot and violin charts for the quantitative data type. This was deemed a minimal data set to assess the effectiveness of each chart and the assumption that a chart failing at this minimal level would likely be unsuitable for a larger quantity of data.

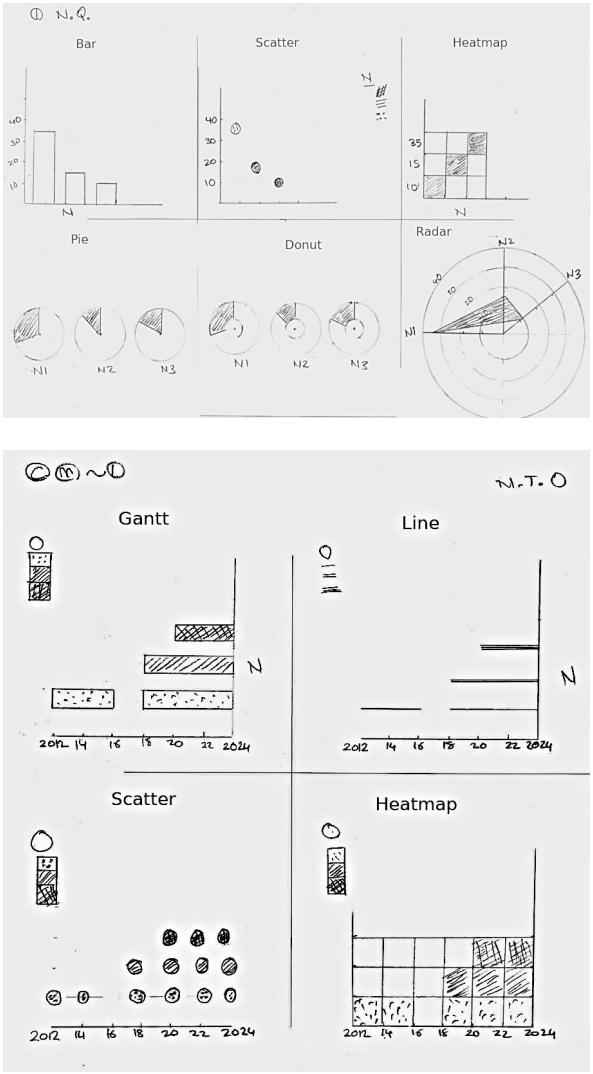

Figure 2: The figure shows six chart options on the top for combinations of nominal and quantitative (NQ) data and four (bottom) for combinations of nominal, temporal and ordinal (NTO) options.

The procedure for evaluating the design choices involved presenting the charts to participants using Microsoft Forms. The form contained two questions about the participant's self-identified "profession" and "years in the current profession". There were also six questions; one for each combination of data types. Each of these questions included an image with the design choices for the combination of data such as those in Figure 2. For each of the six questions, the participants were asked this question: "To what extent do you agree with the statement: 'The chart effectively communicates the information it contains'?". There was also this instruction to answer the question "Select 'Agree' if you think the charts effectively communicate the data, 'Neutral' if you are uncertain or find some charts more effective than others, and 'Disagree' if you think the chart fails to communicate the data effectively." The responses were collected using Microsoft Forms for analysis which will be discussed next.

4.2.2 Results

A total of 18 participants returned the questionnaire. Individual results for each of the six combinations of data types showed clear winners, such as histograms having no "disagree" answers when looking at quantitative data, and likewise, the Gantt chart having no "disagree" answers in the category of nominal and temporal data types as shown in Figure 3. Notably, violin plots received few (3 out of 18) "agree" answers in the quantitative data combination, and donut (doughnut) chart had no "agree" answers in the nominal and temporal data type combination.

Looking at each question and starting with question three in Figure 3, the histogram had the highest "agree" responses given the small data set used in the study. In question four, when looking at the choices for showing a list of names (nominal only), the difference between "agree" answers for "List (Table" and a pie chart is small (66% versus 61%). Question five included the design choices for nominal and quantitative data using seven different chart types. Bar charts had the highest "agree" answers. Question six was showing nominal and temporal data, only Gantt and line charts had no "disagree" answers, while donut chart was the only choice with no "agree" answers. Question seven included four design choices for nominal, temporal and ordinal data, only scatter plots received less than 50% "agree" answers. Question eight used the same four choices from question seven to visualise nominal, temporal and quantitative data. All but Gantt charts received less than 50% "agree" responses.

4.2.3 Discussion

The evaluation aimed to find which charts were effective for representing the simple data set. A starting point was to include each chart where at least 50% of participants agreed that the chart was effective for a given combination of data types. However, there were exceptions.

The choices for quantitative data comprised a histogram, box plot and dot plot. However, box plots only show unimodal distributions and histograms may hide sophistcated data distributions when the default number of bins (e.g., 10) is used. By contrast, violin plots are designed for multimodal distributions and

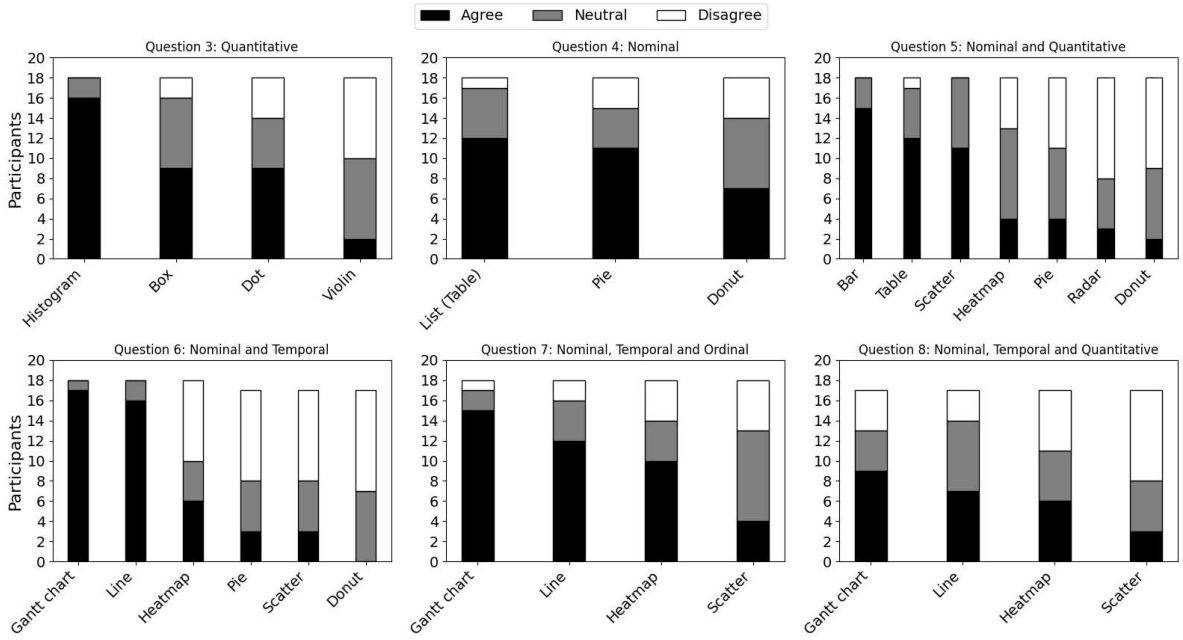

Figure 3: The individual results of the charts design choice evaluation. Each chart is labelled with the question number within the questionnaire and the category of data types for which the design choices were shown to the participants.

those often occur in health data (Ruddle et al., 2024). It is also worth noting that pie and donut charts are essentially the same as each other, except that a donut has a hole in the middle.

5 Implementation

The final step of our design study was to implement and illustrate the final charts using primary care data about patients' conditions from Clinical Practice Research Datalink (CPRD) (Herrett et al., 2015) in a Python package.

The DynAIRx project focuses on conditions from the Quality and Outcomes Framework (QOF) (Roland and Guthrie, 2016) and, in a process that included reviews by clinicians, mapped 4724 SNOMED Clinical Term codes to 25 conditions that were based on QOF conditions (Aslam et al., 2025).

We filtered the CPRD observation table so that it only contained records for the SNOMED codes in the mapping, and then extracted the data for the 2,165,067 patients who had four or more of the conditions. Descriptive statistics were calculated for the number of unique SNOMED code dates, the average gap between those dates and the time-span (max date - min date) for each patient/condition combination. To illustrate the charts with histories that may be considered typical of patients with multiple long-term conditions, we selected the patients who had 4 – 6 conditions and had at least one condition that was in the inter-quartile range for the number of unique dates, average gap and time-span (see Figure 4).

The names of conditions are suitable for answering questions such as "Has anybody ever mentioned something called COPD to you?" and for that nominal data the two visualization techniques on the shortlist were a list (table) and a pie chart. Lists are widely used in current GP systems (e.g., EMIS). A guideline for pie charts is that they should not have more than six wedges (Hardin et al., 2014), which is dictated by the number of colours or shapes people can easily distinguish (Munzner, 2014). That means that a pie chart could be used to indicate conditions for a typical multiple comorbidity patient. However, a pie chart would not be appropriate for patients with the most complex health conditions (half a million patients in our dataset had 7 or more conditions).

Other questions involved nominal and temporal data, e.g., "Did that [diabetes] start later on". The two shortlisted visualization techniques were Gantt and line charts, which both display continuous data. Each row in a medical record indicates an event (e.g., a diagnosis, a prescription or a test result) that occurred on a specific date. The SNOMED codes mapped to each condition (Aslam et al., 2025) provide ongoing indicators that a patient still has (or

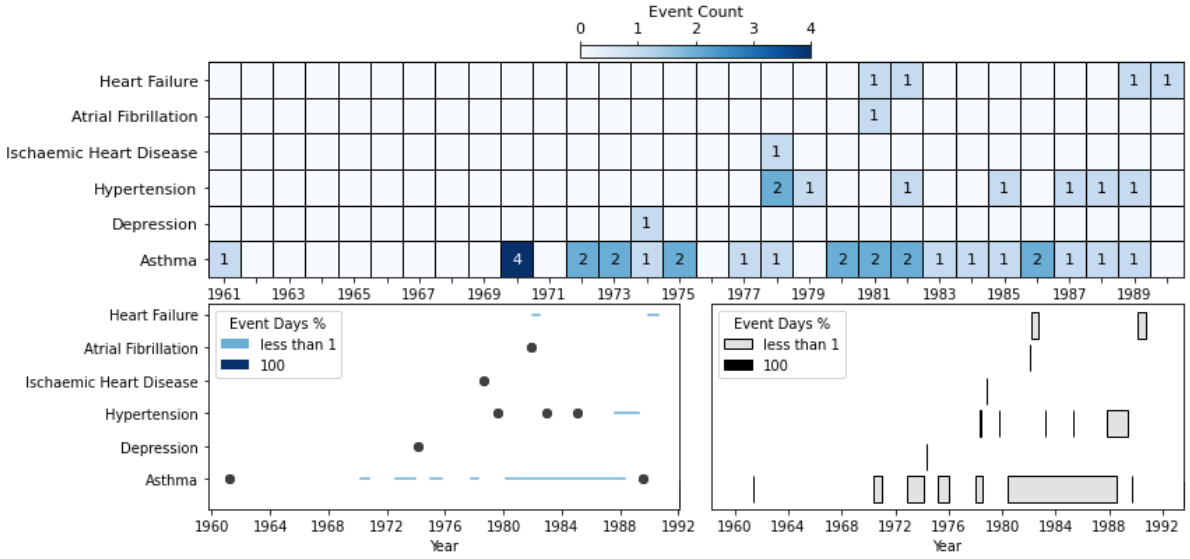

Figure 4: A Gantt chart, line chart and heatmap showing visual histories of a patient's conditions and the event frequency within each time period.

is being investigated for) a given condition, so our approach is to graphically join condition events that occur within a certain threshold (e.g., a year) of the previous one into a bar (Gantt chart) or a line. Events that do not satisfy the threshold are shown as discrete vertical lines (Gantt chart) or dots (line charts), as illustrated in Figure 4.

6 Conclusions and future work

The present work started with reviewing techniques to visualise EHR. Using four mock SMRs the domain problem was characterised and data abstraction was generated to carry out the chart-type design space exploration described in the previous sections. Through exploration of the design space, a list of charts was selected and evaluated. The designs were then implemented in a Python package to generate the same visualizations using CPRD data for 2.1 million real patients. We anticipate that the Python package can be usable for other datasets similar to the CPRD data types used in this study.

The contributions of the present work are based on the three main contributions outlined in the design study method paper (Sedlmair et al., 2012): problem characterisation, that is, analysis and data abstraction using mock SMRs, the design study to present candidate chart types and a Python package-based implementation.

The present work enables us and other visualization designers looking at generating "integrated" dashboards to use the chart types in this study and the Python package implementation to visualise complete patient records to assist clinicians conducting medication review such as SMRs.

ACKNOWLEDGEMENTS

This study is funded by the National Institute for Health Research (NIHR) under its Programme Artificial Intelligence for Multiple and Long-Term Conditions (NIHR 203986). The views expressed are those of the author(s) and not necessarily those of the NIHR or the Department of Health and Social Care. The authors acknowledge the use of the LASER platform, and assistance provided by the Data Analytics Team in Leeds Institute for Data Analytics (University of Leeds). M. Gabbay is part funded by the National Institute of Health and Care Research Applied Research Collaboration North West Coast.

Data availability The research used a pseudonymized dataset that was provided under a CPRD Study Dataset Agreement (Protocol 23_002781). Due to data governance restrictions, the dataset cannot be made openly available.

REFERENCES

Abuzour, A., Wilson, S. and Woodall, A. et al. (2023). How might dynamic artificial intelligence (DynAIRx) be used to support prescribing to ensure efficient medication reviews? *Annals of Family Medicine*, 21(Suppl 3):4823.

Aslam, A., Walker, L. and Abaho, M. et al. (2025). An automation framework for clinical codelist development validated with UK data from patients with multiple long-term conditions. *BMC Medical Research Methodology*, 25(1):1–17.

Calzoni, L., Clermont, G. and Cooper, G. F. et al. (2020). Graphical presentations of clinical data in a learning electronic medical record. *Applied Clinical Informatics*, 11(04):680–691.

Daley, K., Richardson, J. and James, I. et al. (2013). Clinical dashboard: Use in older adult mental health wards. *The Psychiatrist*, 37(3):85–88.

Dowding, D., Randell, R. and Gardner, P. et al. (2015). Dashboards for improving patient care: Review of the literature. *International Journal of Medical Informatics*, 84(2):87–100.

DVS (2021). State of the industry report. `https://www.datavisualizationsociety.org/survey`. Accessed: 2023-11-21.

DynAIRx Project (2025). dynairxvis: Visualisation tools for structured medication review. `https://github.com/DynAIRx/dynairxvis`. Accessed: 2025-06-25.

Elshehaly, M., Alvarado, N. and McVey, L. et al. (2018). From taxonomy to requirements: A task space partitioning approach. In *Proceedings of the IEEE VIS Workshop on Evaluation and Beyond–Methodological Approaches for Visualization (BELIV)*, pages 19–27. IEEE.

EMIS Health (2023). EMIS Health - transforming healthcare through innovation. Accessed: 2023-04-30.

Hardin, M., Hom, D. and Perez, R. et al. (2014). *Which chart or graph is right for you?* Tableau.

Herrett, E., Gallagher, A. M. and Bhaskaran, K. et al. (2015). Data resource profile: Clinical practice research datalink (CPRD). *International Journal of Epidemiology*, 44(3):827–836.

Hripcsak, G. (1994). Writing arden syntax medical logic modules. *Computers in Biology and Medicine*, 24(5):331–363.

Kheterpal, S., Shanks, A. and Tremper, K. K. (2018). Impact of a novel multiparameter decision support system on intraoperative processes of care and postoperative outcomes. *Anesthesiology*, 128(2):272–282.

Kontopantelis, E., Stevens, R. J. and Helms, P. J. et al. (2018). Spatial distribution of clinical computer systems in primary care in England in 2016 and implications for primary care electronic medical record databases: A cross-sectional population study. *BMJ Open*, 8(2):e020738.

Koopman, R. J., Kochendorfer, K. M. and Moore, J. L. et al. (2011). A diabetes dashboard and physician efficiency and accuracy in accessing data needed for high-quality diabetes care. *The Annals of Family Medicine*, 9(5):398–405.

Kraus, S., Drescher, C. and Sedlmayr, M. et al. (2018). Using arden syntax for the creation of a multi-patient surveillance dashboard. *Artificial Intelligence in Medicine*, 92:88–94.

Madden, M., Mills, T. and Atkin, K. et al. (2022). Early implementation of the structured medication review in england: a qualitative study. *British Journal of General Practice*, 72(722):e641–e648.

Munzner, T. (2014). *Visualization Analysis and Design*. A K Peters/CRC Press, Boca Raton, FL, 1st edition.

NHS England (2021). Network contract directed enhanced service structured medication reviews and medicines optimisation: Guidance 2021/22.

OpenAI (2023). Chatgpt-4. `https://www.openai.com`. Accessed: 2023-10-21.

Roland, M. and Guthrie, B. (2016). Quality and outcomes framework: What have we learnt? *BMJ*, 354.

Ruddle, R. A., Cheshire, J. and Fernstad, S. J. F. (2024). Tasks and visualizations used for data profiling: A survey and interview study. *IEEE Transactions on Visualization and Computer Graphics*, 3:3400–3412.

Satyanarayan, A., Russell, R. and Hoffswell, J. et al. (2015). Reactive Vega: A streaming dataflow architecture for declarative interactive visualization. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):659–668.

Schulz, E. B., Phillips, F. and Waterbright, S. (2020). Case-mix adjusted postanaesthesia care unit length of stay and business intelligence dashboards for feedback to anaesthetists. *British Journal of Anaesthesia*, 125(6):1079–1087.

Sedlmair, M., Meyer, M. and Munzner, T. (2012). Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440.

Seitamaa-Hakkarainen, P. and Hakkarainen, K. (2000). Visualization and sketching in the design process. *The Design Journal*, 3(1):3–14.

Snyder, H. (2019). Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104:333–339.

Spence, R. (2001). *Information Visualization*. Springer, London, 1st edition.

Stinson, L. J., McMurray, J. H. and Eaton, C. B. (2012). E-health blood pressure control program. *The Journal of Medical Practice Mnagement: MPM*, 28(2):91.

Wang, Q. and Laramee, R. S. (2022). EHR STAR: the state-of-the-art in interactive EHR visualization. In *Computer Graphics Forum*, volume 41, pages 69–105. Wiley Online Library.