Proceedings of the 2022 Pre-ICIS SIGDSA Symposium

Special Interest Group on Decision Support and Analytics (SIGDSA)

Winter 12-10-2022

A Decision Support System for Managing Uncertainty in the Delivery of Palliative Care in the Community

Ciara Heavin

Helen Leahy

Fiona Kiely

S. Armagan Tarim

A Decision Support System for Managing Uncertainty in the Delivery of Palliative Care in the Community

Research-in-Progress

Ciara Heavin
University College Cork, Ireland.
c.heavin@ucc.ie

Helen Leahy
Marymount Hospice, Cork, Ireland.
hleahy@marymount.ie

Fiona Kiely
Marymount Hospice, Cork, Ireland.
drfkiely@marymount.ie

Armagan Tarim
University College Cork, Ireland.
armagan.tarim@ucc.ie

Abstract

This research-in-progress article addresses Clinical Nurse Specialist (CNS) daily patient prioritisation, scheduling, and routing problem arising in specialised palliative care provided to patients in the community setting. This problem is quite challenging due to the level of uncertainty across the most important elements of these health service operations. Uncertainty inherent in service and travel times, patient behaviour, and patients' healthcare complexity characterise the general theme of uncertainty. Existing research investigates the home healthcare (HHC) resource optimisation problem, some have specifically focused on HHC within the context of palliative care in the community. However, these studies use a selection of metaheuristic approaches and address the HHC resource optimisation problem under the limiting assumption that service operation parameters are deterministic. To solve this problem, we use novel methods to generate data and combine predictive and prescriptive analytics approaches to deliver an AI-driven decision support system (DSS) used by healthcare staff and patients to support the allocation and delivery of specialised palliative healthcare in the community. This new dataset and these analytics capabilities will provide healthcare services with a new and sensitive means to make evidence-based decisions about the delivery of these important healthcare services.

Keywords

Decision support system, Predictive analytics, prescriptive analytics, Artificial Intelligence (AI), Operations Research (OR), Palliative Care (PC), healthcare

Introduction

It is estimated that by 2050 the global over 65s population will have doubled to 1.5 billion. As people live longer, the relative size of the labour force will shrink meaning there will be less healthcare budget per older person. In line with this, all countries face a growing need for palliative care services. Palliative care "improves the quality of life of patients and their families facing the problems associated with life-threatening illness, through the prevention and relief of suffering by means of early identification and impeccable assessment and treatment of pain and other problems, physical, psychosocial, and spiritual" (May et al., 2019). By 2046 in Ireland alone, it is expected there will be an 84% increase in the number of people requiring palliative care (May et al., 2019). Coupled with a heavy disease burden, there is an immediate need to develop the healthcare workforce and grow palliative care service capacity (Sleeman et al., 2019). Investment in palliative care can relieve suffering for patients and families and save money for healthcare systems and society more broadly (Chalkidou et al., 2014).

Existing research has investigated the HHC resource optimisation problem (Carello and Lanzarone, 2014; Heching et al., 2019). Researchers have specifically focused on HHC within the context of palliative care in the community (Heching et al., 2019). These studies rely on diverse metaheuristic approaches including tabu search (intelligent search that support algorithm optimisation (Hertz and Lahrichi, 2009; Rest and Hirsch, 2016), pattern or column generation (Allaoua et al., 2016; Cappanera and Scutellà, 2014), variable neighborhood search (Mankowska et al., 2018; Trautsamwieser and Hirsch, 2011)., variable neighborhood search combined with scatter search and other heuristics (Hiermann et al., 2015), constraint programming combined with heuristics (Nickel et al., 2012; Rendl et al., 2012), an inexact Benders method (Cire and Hooker, 2012), implementing a separate solution comprising rostering and scheduling components (Yalçındağ et al., 2016) to interrogate this problem. The aforementioned work addresses the HHC resource optimisation problem under the limiting assumption that service operation parameters are known, which is generally not the case in healthcare.

Given the complexity of palliative care needs and the scarcity of available specialised resources, uncertainty is among the most critical elements of today's service operations and requires significant attention. Uncertainty may be characterised by unpredictable availability of service which has heightened against the Covid-19 backdrop, patient service time, and nurse travel times where patients are located considerable distances away from the specialist palliative care healthcare service provider. Further, patient behaviour (i.e., the availability and preferences of patients receiving care), and patients' healthcare complexity (i.e., patient stability may fluctuate depending on medical and social factors) highlight the extent to which uncertainty prevails in the delivery of specialist care in a community setting. While this uncertainty is a complicating factor for planners and schedulers, planning and scheduling systems rarely support non-deterministic views of data. It is known that our inability to effectively address uncertainties in operations planning is an increasingly critical limitation and a significant barrier to the use of research results in practice (Wu et al., 1999).

This research-in-progress study provides a high-level view of the steps taken towards the design and development of the CommPAL solution. Using Artificial Intelligence (AI) and Machine Learning (ML), we present an innovative approach to manage uncertainty and to respond rapidly and reliably to changing conditions in the palliative care setting.

This research in progress article is structured as follows: the next section provides a brief outline of the problem identification phase which was completed in collaboration with a large regional specialised palliative care service in the south of Ireland. Next, the main components of the CommPAL solution are presented and the methods proposed to design the triaging and routing and rostering models are outlined. Conclusions are presented.

Study Context: CommPAL Project

This research commenced in 2020 when we undertook a market investigation of the state of the art in AI in the home healthcare and palliative care domains (Kupper et al., 2020). Our investigation revealed that there is no existing solution that enables a data-based approach to the allocation of finite healthcare resources, while supporting the predictive triaging of patient care plan stability in palliative community outreach services. Further, two workshops (including 33 participants) and 11 face-to-face interviews were conducted with key stakeholders with expertise in the specialist palliative care domain. These included the CEO, Clinical Nurses Specialists (CNS), Consultants, Pastoral and Social Care workers, a local General Practitioner and Occupational Therapists. Using De Bono's Six Thinking Hats together with a persona building approach enabled us to gather rich data capturing stakeholders' experiences relating to the delivery of a specialist palliative care service. A thematic analysis of this data allowed us to identify key insights about the delivery of healthcare in the community: how the multidisciplinary team feel about what they do, the strengths of the service and the aspects of the service that could be improved, the key stakeholders, their tasks, pains, opportunities, and challenges. Healthcare professionals expressed sentiments of feeling overburdened by the amount of work and the limited resources available. As one healthcare professional noted, decision making around service delivery is increasingly inconsistent. Reducing existing cognitive overload and subsequent *"decision fatigue"* would improve staff well-being, as well as the quality of the service. This early work provided the foundation on which a decision support system (DSS), namely CommPAL, is designed, as outlined in the next section.

CommPAL Decision Support System

CommPAL involves the design and development of a DSS to support the delivery of specialised palliative healthcare in the community. The aim is to 1) maximise the number of patients cared for in a fair way, 2) better support CNSs by providing them with an optimum and transparent plan to undertake their roles, 3) better understand the changing health needs of patients, 4) get the right information to healthcare professionals at the right time to support decision making to reduce uncertainty, 5) to improve the wellbeing of people with life-limiting illnesses and their families, and 6) optimise the routes taken by CNS with the aim of delivering a more sustainable healthcare service. The solution will combine predictive and prescriptive analytics.

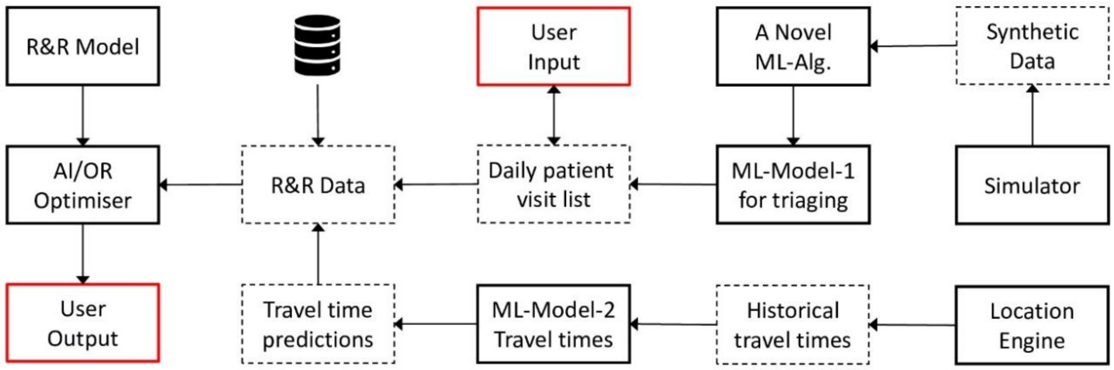

Figure 1. Conceptual view of proposed DSS

Figure 1 illustrates the main components of the proposed CommPal solution:

- The "simulator", which is designed in collaboration with expert healthcare decision makers, is used to generate high-fidelity synthetic data for palliative care.
- The generated realistic data points (i.e., feature vectors) for "synthetic patients" are class labelled, where the categories represent the levels of urgency for patients, by expert healthcare decision makers.
- The labelled synthetic data is used to train a predictive model by a novel "machine learning algorithm for triaging".
- The resultant machine learning model, "ML-Model-1", is deployed to process the most up-to-date patient data and to generate a daily patient visit list.
- The "daily patient visit list", which essentially comprises an ordered list of patients with respect to a set of social and healthcare measures, is presented to the healthcare decision maker to review and finalise the evaluation. At the end of this step, the CNS daily patient visit list is generated and recorded as part of the routing-and-rostering (R&R) data to be used for creating an instance of the R&R optimisation model.
- In parallel to these steps, a "location engine" that is designed specifically to collect travel times and other relevant data in a given city or a service district is run.
- The collected "historical travel times" and other data are input into another machine learning model ("ML-Model-2") to predict the travel times for a given day.
- The outputs of "ML-Model-2" are the "travel time predictions", which are not point forecasts, but in the form of probability distribution functions. These distributions too are recorded as part of the R&R data.
- The collated "R&R data" comprises information related to nurses and patients, as well as the CNS daily patient visit list and travel time distributions.
- An instance of the "R&R model", which is essentially a hybrid AI/OR optimisation model, is created using the "R&R data" for the given day.
- The optimal (or a near-optimal) solution to the instance is obtained by using an AI/OR model optimiser. This solution fundamentally corresponds to a detailed recommended daily patient visit plan involving rostering and routing decisions. The daily plan presented to the CNS can be revised in the final step of the decision support process.

In brief, the aim of the predictive analytics component is to assess patient care needs and the stability of a patient's condition and to help to triage patients. This will include patients self-reported input (patient reported outcome measures) about their own wellbeing using a mobile device. Predictive analytics will also be used to determine the distributions of travel times and durations of patient visits. The outputs from the predictive analytics components will be used as inputs to the prescriptive analytics component, whose aim is to assist with the decision making by allocating resources in a fair, and a transparent way.

Method

Phase 1: Patient Triage

We designed and developed a data model for palliative healthcare services in the community. Using the newly implemented Palliative Assessment and Clinical Response Palliative Care outcome collaboration (PCOC) and other socio-demographic features, we present a robust data model that allows for the capture of data from patients receiving care from a CNS. Predictive analytics approaches and ML techniques will be used to analyse patient data to predict patient healthcare needs in the future i.e., prioritise patients based on their current level of stability over past visits. In consultation with healthcare domain experts, in total, twenty-one different patient attributes are identified as the main determinants of a patient's stability level. The attribute set includes:

> A: Age; B: Degree of social support; C: Lives alone; D: Residential care; E: Diagnosis; F: Multimorbidity (> 2 chronic illnesses).
> G: Evidence of POD; H: Ongoing disease modifying treatments.
> I: Sleeping; J: Appetite; K: Nausea; L: Bowels; M: Breathing; N: Fatigue; O: Family/Carer; P: Pain (PSS); Q: Psych/Spiritual.
> R: CSCI required; S: RUG (total).
> T: AKPS; U: Urgent need for crisis event planning.

In the development of machine learning models, a synthetic dataset is resorted due to the non-existence of historical data comprising all the above attributes. The synthetic dataset is used to create artificial healthcare records that cannot be attributed to any real individual, but yet are expected to be plausible in the sense that all attributes are coherent with respect to each other. For example, a synthetic data generation method should not result in a female patient with a record of prostate cancer.

Healthcare records are often multivariate categorical data posing multiple modelling challenges, particularly in high dimensions. As a naïve approach to synthetic data generation, one might consider creating a record simply by sampling from the empirical marginal distribution of each attribute. However, this approach does not capture statistical dependencies across attributes, and therefore the randomly generated records may fail to capture the underlying characteristics of the population (Goncalves et al., 2020).

In this paper, to capture the correct distributions and dependencies between attributes, based upon domain expert knowledge, a custom-built hand-coded generative model of data is developed. The resultant model is then used to generate synthetic data via sampling.

The causal dependencies between the attributes of the generative data model are often represented as a *structural causal model* (SCM) given in the form of a graph, G, that provides an intuitive visualisation by representing attributes as nodes, $V = \{1, \dots, n\}$, and relationships between attributes as edges in this graph. Often, these edges are directed edges, D, indicating the direction of causality.

Bayesian Networks (BN) are one of the most widely used SCMs and used to represent the dependency relationships of attributes and their joint distributions efficiently in a factorised way. Each BN describes these dependencies as a *directed acyclic graph* (DAG) and the decomposition of the joint probability distributions as a product of conditional probability distributions constrained by the structure of the DAG.

A DAG, $G = (V, D)$, together with a collection of conditional probability distributions that are provided by domain experts, is a directed graph in which there is no directed cycle. Each node is associated with a conditional probability distribution of the corresponding attribute given its parents in the DAG. Node $j \in G$

is a parent of a node $i \in G$ if there is a directed edge from j to i in the graph G. Then a joint probability distribution factorises over the directed graph G if,

$$p(x_1, \dots, x_n) = \prod_{i \in V} p\left(x_i \middle| x_{\text{parents}(i)}\right).$$

where $\vec{x}$ denotes the attribute vector. In other words, given a distribution $p(x_1, \dots, x_n)$ and a graph G, p factorises over G if p can be expressed as a product of each attribute's $p\left(x_i \middle| x_{\text{parents}(i)}\right)$.

The synthetic data generated using the defined DAG capture correct (potentially non-linear and multivariate) relationships and distributions that are apparent in the real data sets, while also preserving patient privacy and avoiding the risks of individual identification.

A matrix-DAG representation of the SCM of patient stability used in this work is presented in Figure 2.

	G	H	I	J	K	L	M	N	O	P	Q	R	S	T
A	●	●	●			●				●				●
B		●	●							●	●			
C			●							●	●			
D		●	●							●	●	●		
E	●	●	●			●	●			●	●	●		●
F	●	●	●			●	●			●		●		●
G			●			●	●				●	●		●
H			●			●	●				●	●		●
I				●	●	●	●	●	●	●	●	●		●
J					●	●	●	●	●	●	●	●		●
K						●	●	●	●	●	●	●		●
L							●	●	●	●	●	●		●
M								●	●	●	●	●		●
N									●	●	●	●		●
O										●	●	●		●
P											●	●		●
Q												●		●
R													●	●
S														●

Figure 2. A matrix-DAG representation of the SCM of patient stability

A synthetic data model is built based upon the relationships between the attributes depicted in Figure 2 and a simulator is developed using this data model to generate a set of random “plausible” patients. The attribute vectors of these randomly generated patients are presented to the domain experts by means of a scatter plot matrix to help them easily visualise bivariate relationships between combinations of attributes and confirm the validity of the underlying distributions in the light of their professional experiences.

In the next phase of our research, we intend to have these randomly generated attribute vectors labelled by domain experts, and ML algorithms will be developed to classify patients into "needs urgent visit" and "not-urgent" categories. This work is going.

Phase 2: Routing and Rostering (Patients in the community and CNS)

The routing and rostering (R&R) model focuses on a given day to schedule healthcare professional visits to patients with any medical condition that requires specialist palliative care services in the community. The sequence that a healthcare professional visits patients within a set geographical location to ensure optimum cost and greenhouse gas emissions is based on:

- patient triage results from the triage module,
- healthcare professional location,
- patient location,
- patient availability,
- healthcare professional availability,
- hospital inputted healthcare professional care specialty / skillset,
- GHG emissions with regard to healthcare professional transport mode,
- medical equipment availability.

The output provides a recommendation to an expert planner on a given day. The proposed model guarantees that the conduct of planning activities is transparent and that the results are fair and equitable.

Discussion

Solutions to the challenge of improving operational efficiency and planning of activities in the optimal fashion must address, or at a minimum be compatible with, important aspects of the healthcare environment in which service is provided. As the recent COVID crisis showed once again, uncertainty is among the most critical elements of this environment and requires significant attention by researchers and DSS developers. Random behaviours like uncertain process times, lack of essential information, incorrect data, and vague or incomplete definitions exemplify the general theme of uncertainty in healthcare provision in the community. The importance of modelling uncertainty in such decision support systems is growing every day. Innovative original approaches to manage uncertainty and to respond rapidly and reliably to changing situations are crucial. While uncertainty is seen as a complicating factor for planners and schedulers, planning and scheduling systems rarely support non-deterministic views of data. As a visible example in the business analytics context, Enterprise Resource Planning (ERP) and other corporate information and planning systems seldom incorporate mechanisms to consider variance or other information on randomness. It is known that our inability to effectively address uncertainties in operations planning is an increasingly critical limitation and a significant barrier to the use of research results in practice (Wu et al., 1999).

To address the aforementioned issues and generate robust R&R plans, uncertainty inherent in patient visit durations and travel times will be explicitly considered in the modelling process. The resultant models are expected to be in the form of Stochastic Programs/Stochastic Constraint Programs (SP/SCP). These mathematical models are non-stationary stochastic combinatorial optimisation models, which are notoriously difficult to solve to optimality. Therefore, to obtain optimal or near-optimal solutions to the R&R problem, off-the-shelf commercial solvers such as Cplex/Gurobi/XPress-MP, dedicated solution algorithms developed to solve the joint R&R problem, and heuristic methods will be tested.

Conclusion

With limited resources and increasing demand, largely attributed to increased life expectancy of patients and increasingly complex patient comorbidities (e.g., cancer and dementia), it is critical that the proposed solution incorporates the needs of those living longer-term at home with serious illness. This research offers a significant opportunity to transform the delivery of an important service to those patients and their families most in need of care in the community. With the proposed CommPAL solution, a unique and important opportunity exists to apply analytics techniques to combine triaging patients using palliative care clinical guidelines with the concomitant planning of the delivery of specialist palliative care and rostering.

Our proposed solution leverages novel approaches such as: stochastic optimisation using AI/OR to address the high-level of uncertainty of the decision problem, the combination of predictive and prescriptive analytics and fairness/ethics in AI to deliver a data-driven solution to support healthcare professionals' decision-making needs. The next step in this research involves the design and testing of the models outlined in this work. Beyond this, our aim is to identify additional clinical test sites to assess the validity and reliability of the proposed solution.

Acknowledgements

This work is supported by Enterprise Ireland Grant no. CF-2021-1676-I.

References

Allaoua, H., Borne, S., Létocart, L., & Calvo, R. W. (2013). A matheuristic approach for solving a home health care problem. Electronic Notes in Discrete Mathematics, 41, 471-478.

Carello, G., & Lanzarone, E. (2014). A cardinality-constrained robust model for the assignment problem in home care services. European Journal of Operational Research, 236(2), 748-762.

Chalkidou K, Marquez P, Dhillon PK, et al. (2014). Evidence-informed frameworks for cost-effective cancer care and prevention in low, middle, and high-income countries. Lancet Oncol; 15: e119–31

Cire A., and Hooker J.N., (2012). A heuristic logic-based Benders method for the home health care problem, presented at Matheuristics, Angra dos Reis, Brazil.

Goncalves, A., Ray, P., Soper, B., Stevens, J., Coyle, L., & Sales, A. P. (2020). Generation and evaluation of synthetic patient data. BMC medical research methodology, 20(1), 1-40.

Heching, A., Hooker, J. N., & Kimura, R. (2019). A logic-based benders approach to home healthcare delivery. Transportation Science, 53(2), 510-522.

Hertz, A., & Lahrichi, N. (2009). A patient assignment algorithm for home care services. Journal of the Operational Research Society, 60(4), 481-495.

Hiermann G, Prandtstetter M, Rendl A, Puchinger J, Raidl GR (2015) Metaheuristics for solving a multimodal home-healthcare scheduling problem. Cent Eur J Oper Res 23(1):89–113.

Kupper, M., Tarim, S. A., Heavin, C., and Kiely, F. (2020) "Towards Transformative Analytics for Palliative Care". AMCIS 2020 Proceedings. 13. https://aisel.aisnet.org/amcis2020/adv_info_systems_research/adv_info_systems_research/13

May, P., Johnston, B. M., Normand, C., Higginson, I. J., Kenny, R. A., & Ryan, K. (2019). Population-based palliative care planning in Ireland: how many people will live and die with serious illness to 2046?. HRB open research, 2, 35.https://doi.org/10.12688/hrbopenres.12975.1

Mankowska, D. S., Meisel, F., & Bierwirth, C. (2014). The home health care routing and scheduling problem with interdependent services. Health care management science, 17(1), 15-30.

Nickel, S., Schröder, M., & Steeg, J. (2012). Mid-term and short-term planning support for home health care services. European Journal of Operational Research, 219(3), 574-587.

Rendl, A., Prandtstetter M., Hiermann G., Puchinger J., Raidl G. (2012). Hybrid heuristics for multimodal homecare scheduling. Beldiceanu N, Jussien N, Pinson E, eds., CPAIOR Proceedings, volume 7298 of Lecture Notes in Computer Science, 339{355 (Springer).

Rest, K. D., & Hirsch, P. (2016). Daily scheduling of home health care services using time-dependent public transport. Flexible Services and Manufacturing Journal, 28(3), 495-525.

Sleeman, K. E., de Brito, M., Etkind, S., Nkhoma, K., Guo, P., Higginson, I. J., ... & Harding, R. (2019). The escalating global burden of serious health-related suffering: projections to 2060 by world regions, age groups, and health conditions. The Lancet Global Health, 7(7), e883-e892.

Wu, S. D., Roundy, R. O., Storer, R. H., Martin-Vega, L. A. (1999). Manufacturing logistics research: taxonomy and directions, Cornell University OR/IE, Technical Report: 1254.

Yalçındağ S, Cappanera P, Scutellà MG, Şahin E, Matta A (2016) Pattern-based decompositions for human resource planning in home health care services. Comput Oper Res 73(12).