Marketing Intelligence and Planning xx (2023) 1–16

Maximizing Customers' Lifetime Value using Limited Marketing Resources

Mage Marmol[a], Anita Goyal[b], Pedro J. Copado[c,a], Javier Panadero[c,a] and Angel A. Juan[c,a]

[a]*Management Science Dept., Euncet Business School, Terrassa, Spain*
[b]*Indian Institute of Management Lucknow, Lucknow, India*
[c]*IN3 – Computer Science Dept., Universitat Oberta de Catalunya, Barcelona, Spain*
E-mail: mage.marmol@euncet.es [M. Marmol]; anita.g@iiml.ac.in [A. Goyal]; pcopadom@uoc.edu [P. Copado]; jpanaderom@uoc.edu [J. Panadero]; ajuanp@uoc.edu [A. A. Juan]

Received DD MMMM YYYY; received in revised form DD MMMM YYYY; accepted DD MMMM YYYY

Abstract

The profitability of customers for a business enterprise can be estimated by the so-called customer lifetime value (CLV). One specific goal for many enterprises consists in maximizing the aggregated CLV associated with its set of customers. To achieve this goal, a company can employ commercial actions (e.g., segmented or even fully personalized marketing campaigns), which are usually expensive. There is an economic trade-off between the investment in marketing actions and the aggregated CLV that can be achieved, i.e.: the higher the marketing budget, the higher the aggregated CLV that can be achieved. Hence, after reviewing the existing literature on optimization of the CLV, this paper proposes an original model to optimize the aggregated CLV subject to an available marketing budget. A solution to this problem is a selection of commercial actions to be deployed. The suggested model is inspired by the well-known uncapacitated facility location problem, where facilities servicing customers represent commercial actions, and the 'distance' from a customer to an action represents how close the action meets the customers needs and interests (hence, the shortest the dinstance, the higher the increase in CLV that can be achieved). The aforementioned concepts are illustrated with a case study example. Finally, a series of computational experiments show the potential of the proposed approach when compared with a standard (non-algorithmic) one.

Keywords: intelligent marketing; customer lifetime value; algorithms in marketing

1. Introduction

In today's competitive world, many managers consider customers' satisfaction and retention as the basis for loyal (long-term) relationships that can provide sustainable incomes over time. Customers' satisfaction alone does not necessarily lead to customers' loyalty to the firm. Actually, some researchers suggest

*Author to whom all correspondence should be addressed (e-mail: mage.marmol@euncet.es).

that customer's satisfaction and loyalty are not always strongly correlated (Ganiyu et al., 2012). For non-loyal customers, satisfaction is mostly based on the current transaction. For loyal customers, satisfaction seems to be related to the accumulated experience (Yi and La, 2004). In principle, satisfied customers are more likely to become loyal ones, who are also more likely to repurchase products or services from an organization. Still, customers' loyalty needs to be always promoted by means of commercial actions, such as marketing campaigns, special offers, etc. (Chang, 2020).

Accordingly, many companies have traditionally used relationship marketing to build up and maintain a base of committed customers (Garepasha et al., 2020). A long-term relationship with customers can generate a competitive advantage for a company. For instance, it can help to disseminate their products and services among potential customers via positive word-of-mouth actions. Therefore, the concept of customer's loyalty is central to relationship marketing (Hennig-Thurau et al., 2001). One of the first sectors where customer loyalty emerged is the airline industry. Ever since, studies on customer expectations, service performance perception, satisfaction, and loyalty have been extended to service organizations like banks, insurance, hotels, etc (Ehigie, 2006; Mokhtar et al., 2018).

As suggested by Zeithaml et al. (2001), targeting loyal customers can make marketing spending more effective. Selecting customers for targeting and determining the level of resources to be allocated can result in a more efficient use of marketing actions (Rust et al., 2011). Moreover, to be competitive, companies spend large budgets on building long-term relationships with their clients through different marketing campaigns, i.e.: promotions, communications in different media, publicity activities, public relations, etc. Thus, managers must decide how to allocate the resources they have to optimize the spending and the customers' retention. Marketing campaigns might influence individual customers' loyalty value (CLV), which can be defined as "the margin of revenue or profit that a customer provides during the time of her loyalty to the company" (Kumar et al., 2004). Therefore, managers should carefully design the marketing actions to be employed (von Mutius and Huchzermeier, 2021). In this paper, we discuss how to optimize the aggregated CLV via an intelligent selection of marketing actions.

Figure 1 illustrates the trade-off between the activation of resources (e.g., customized marketing campaigns, which tend to be costly) and the CLV. The left side of this figure shows a scenario with low resource-activation costs, since only one generic commercial action is deployed. Still, this scenario also shows low improvements in terms of CLV: we miss the potential increase in aggregated CLV that a more customized set of commercial actions could have achieved. For each customer, the associated CLV increase is assumed to be inversely proportional to the distance between the customer and the active commercial action that is closer to her needs and interests. Hence, the right side of Figure 1 shows a scenario with a higher aggregated CLV (since customers are 'closer' to the deployed actions) and also higher activation costs (since more commercial actions are deployed).

Our goal is then to find a configuration of deployed commercial actions that maximizes the total CLV for the firm while respecting a given budget threshold. Notice that we are making the following assumptions: *(i)* the set of commercial actions that could be deployed is given (these actions could be, for instance, a series of segmented campaigns proposed by the marketing department); *(ii)* the cost of deploying each of the possible actions is known (e.g., this cost could be inferred from previous campaigns of similar characteristics); *(iii)* for each customer, the increase in CLV associated with assigning a marketing action to her can be estimated; and *(iv)* for a given mapping of deployed actions, each customer will be assigned to the active action that achieves a higher increase in her CLV. Typically, historical data can be used, in combination with predictive models (e.g., multiple regression ones), to estimate the im-

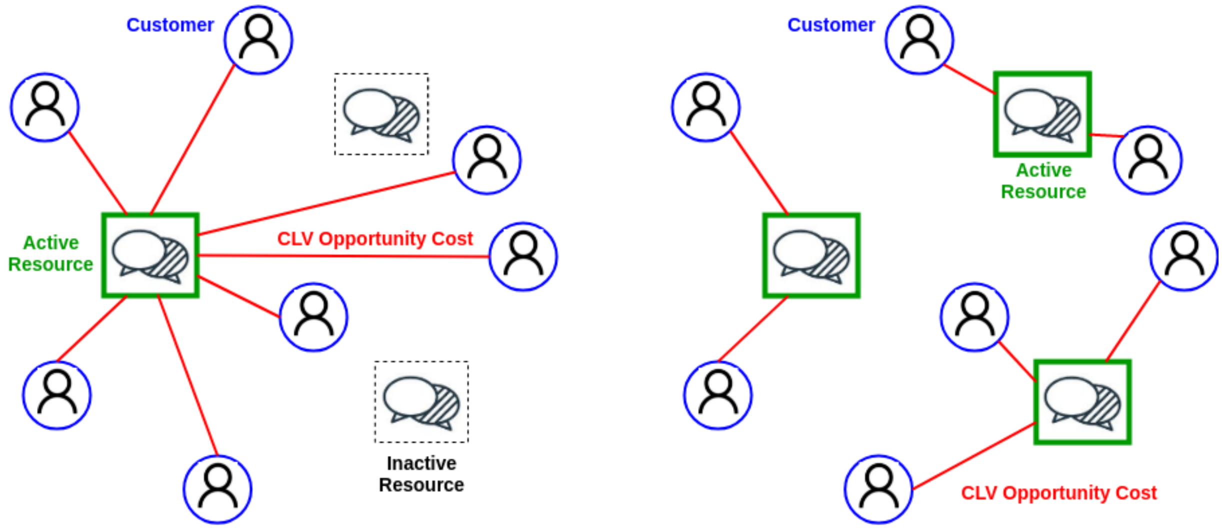

Fig. 1: A visual illustration of the trade-off between active resources and aggregated CLV.

pact on the aggregated CLV generated by the activation of different marketing actions. Notice also that the aggregated CLV can be obtained by estimating the effect of the selected actions on each customer type. Thus, customers are assumed to be classified / assigned into groups or clusters.

The main contributions of this paper are: *(i)* a discussion on the trade-offs between maximizing customer lifetime value and minimizing the cost of deploying expensive commercial actions; *(ii)* a formal model that describes the former trade-off as a variant of the uncapacitated facility location problem; *(iii)* a biased-randomized algorithm (Quintero-Araujo et al., 2017) that is capable of solving the aforementioned optimization problem in short computing times; and *(iv)* a series of benchmark instances that can be used by future researchers to test alternative solving approaches. Biased-randomized algorithms are specially useful in practical applications for different reasons: *(i)* they are relatively simple to implement, flexible, and intuitive for managers; *(ii)* they offer high-quality solutions to complex problems in short computing times (e.g., in a few minutes or even less if parallel-computing techniques are employed); and *(iii)* the setting of their parameters is typically quite simple and can be completed after some quick tests (Dominguez et al., 2014). Hence, while developing our algorithmic approach, we have tried to balance efficiency and simplicity, so the proposed algorithm can be employed in different real-life scenarios. Also, to the best of our knowledge, this is the first work that models CLV management as a facility location problem (Pagès-Bernaus et al., 2019), which introduces a novel view of this topic. Actually, most of the existing studies consider CLV as an *exogenous* variable, i.e.: one that does not depend on the marketing actions. Marketing resources are frequently allocated over the entire customer base rather than across individuals or segments. However, in order to maximize the aggregated CLV, companies need to consider how their customized marketing strategies can affect this variable. By segmenting customers and allocating personalized marketing strategies to every segment, firms can enhance their return in terms of aggregated CLV.

The remaining of this paper is structured as follows: Section 2 reviews the literature associated with the CLV concept and how it can be estimated by employing different methodologies. Section 3 focuses on previous work related to the optimization of the CLV. Section 4 formulates the optimization problem considered in this paper as a variant of the well-known facility location problem. Section 5 proposes a constructive heuristic, and then extend it into a biased-randomized algorithm, which allows us to generate better solutions in short computing times. Section 6 provides a case study to better illustrate the main concepts behind our model with a numerical example. Section 7 describes a series of computational experiments that have been carried out to illustrate the concepts introduced in this paper. Finally, Section 8 summarizes the main contributions and outcomes of this work.

2. Works Related to the Customer's Lifetime Value

Customer's lifetime value has been a popular research topic during the last decades. Some research papers present a comprehensive review of CLV, its benefits, and its computation methods. CLV can be defined as the present value of all the future cash flows attributed to a customer's relationship with an organization (Kumar et al., 2004; Verhoef and Lemon, 2013), and offers the advantage to assess the financial value of each customer –with the aim of identifying the most profitable customers and to nurture long-term relationships (Kumar et al., 2008). According to AboElHamd et al. (2020), CLV is one of the most reliable indicators in direct marketing for measuring the profitability of customers. It is also a core metric in customer relationship management, where a relationship is developed over time based on a series of interactions between parties –customers and service providers. It deals with increasing customer's value and satisfaction, and also with guaranteeing a long-term profitability for the firm. Thus, CLV can be considered at different levels, including: individual, segment, or cohort. CLV can also be expressed as excess revenue over the cost of acquiring, servicing, and retaining a customer (Berger and Nasr, 1998). It involves the study of defection rates, purchases, average spending, and cross-buying behavior. CLV can be useful to improve market segmentation and resource allocation, evaluate competitor firms, customize marketing communication, optimize the timing of product offerings, and determine a firm's market value (Gupta et al., 2004; Kumar et al., 2006).

The CLV is normally computed by looking at the future purchase potential, since customers grow in number with time, thus resulting in higher revenue –and a reduction in cost-to-serve a customer (Reichheld and Kenny, 1990). There are two basic approaches: *(i)* a historical approach that takes into account data about previous purchases; and *(ii)* a predictive approach that applies algorithms to customer's behavior and predicts their future value. Berger and Nasr (1998) had presented general mathematical models to compute the CLV based on a net contribution margin attained by a customer after acquisition. Case examples are presented considering customer retention models (customers either remain or not with the vendor) and customer migration models, i.e., customers switching between vendors (Dwyer, 1997). Some authors propose to combine real-options analysis with the CLV computation to analyze whether to abandon non-profitable customers for meaningful resource allocation. Others consider that CLV without option analysis underestimates the customer value. Factors influencing CLV were identified through Monte-Carlo simulations by Haenlein et al. (2006). In business-to-business relationships, CLV models are suggested in consideration of the four different types of buyer-seller relationships, which are based on a dependence distribution between buyers and sellers. Thus, for instance Roemer (2006) suggested CLV models that are based on either the discounted future cash-flow method or on the discounted cash-flow combined with options value or with possible customer's strategies.

Borle et al. (2008) assessed the CLV of each customer at each purchase occasion using a hierarchical Bayes approach. In the model, they included the purchase timing, purchase amount, and risk of defection for each customer. Romero et al. (2013) used a partially hidden Markov truncated model to obtain more accurate forecasts of future customer's behavior. Their model included purchase incidence decisions, besides considering novel factors such as dynamic purchase patterns, dependence between purchase frequency and monetary value, and customers that can become active again after a few periods of temporary inactivity. Khajvand et al. (2011) used CLV to segment customers in a health and beauty company. They firstly employed the recency-frequency-monetary method. Then, they extended this method with an additional parameter called 'count item'.

Venkatesan and Kumar (2004) have analyzed how CLV is affected with the use of different channels of communication. Models that are able to forecast the customer's purchase frequency (and the associated contribution margin) were suggested to predict the CLV. Rust et al. (2011) also analyzed how CLV-based resource allocation helps to increase revenue without changing the amount of marketing resource investments.

Ekinci et al. (2014a) focus on increasing profits by providing an optimal promotional plan. Their approach was first to have clusters of customers based on their CLVs. Then, they assigned different promotion campaigns to each cluster in order to optimize profit. Hence, these authors considered CLV as an exogenous variable –at least in part–, which helps to find optimal promotion campaigns. They conclude that, with better methods, CLV can be maximized with an appropriate resource-allocation strategy, which is precisely the focus of our study. Ching et al. (2004) classified customers according to their purchases to a given company. Then, the effect of different promotion scenarios on the CLV is considered. In addition, the impact of the budget constraint is also studied. They use a Markov chain to model the customer's stochastic behavior, and they solve the underlying optimization problem by employing either linear programming (infinite horizon) or dynamic programming (finite horizon). Their model, however, does not consider personalized marketing campaigns nor large-scale scenarios with thousands of customers and hundreds of possible customized campaigns, which are central aspects in our study.

3. Related Work on Optimizing the CLV

Any company pursues an optimal allocation of resources, considering both cost and revenue, to maintain profitability and market performance (Kumar, 2006). Such an efficient allocation is required to balance expenditure between acquisition of new customers with a low CLV and retaining old customers with a high CLV (Gupta et al., 2006). With the goal of maximizing the CVL, Venkatesan and Kumar (2004) proposes a genetic algorithm to assess a desired contact interval time for each individual customer. With the same goal as before, Jonker et al. (2004) proposed a customer segmentation process. Fruchter and Sigué (2009) identify, through optimal control theory of linear control systems, that the CLV can be increased by a growth in trust and a decrease in opportunism between the buyer and the seller. Ekinci et al. (2014b) try to increase the CLV by optimizing the type of promotion required for a customer segment. Their approach is based on stochastic dynamic programming, as well as on classification and regression trees. For a scenario with different types of products, these authors presented a promotional plan aimed at maximizing the CLV. Ching et al. (2004) also suggested optimizing the CLV with a stochastic dynamic programming model. In case of promotions with no limits, the linear programming technique is proposed, whereas a dynamic programming approach is considered in a situation with a limited number of promotions. The authors have also included a budget constraint factor with application to real data.

Crowder et al. (2007) have given a framework to optimize the CLV considering some features. The main one is the time duration of the probationary period, i.e., for how long a customer will be monitored to know her profitability. The authors have taken into consideration the variability within individuals to describe a stochastic behavior, as well as the variability among individuals. Blattberg et al. (2009) have discussed conceptual issues towards the effect of various marketing elements, like price, promotion, cross buying, multi-channels, RFM, and others towards CLV maximization. Ekinci et al. (2014b) have presented a review on CLV optimization considering a specific marketing decision, such as promotional budgets, sales-force allocation, product recommendation, customer acquisition, or retention rate.

AboElHamd et al. (2020) have reviewed traditional and dynamic programming models to maximize the CLV, highlighting the advantages and disadvantages of these models. Some basic or traditional methods included in their study are: fuzzy systems, neural networks, system dynamics, and Bayesian decision theory. These are primarily based on CLV indicators like churn rate, retention rate, etc. Their dynamic programming algorithms include Markov decision processes, approximate dynamic programming –which is also indicated as reinforcement learning–, and deep learning. AboElHamd et al. (2020) have also proposed the use of deep learning for CLV maximization.

The existing literature presents an understanding of the CLV with its important place in customer relationship management, as well as models to measure and maximize it. Most studies also conclude that CLV contributes to customer segment development and resource allocation. Despite the amount of literature on the variables taken into account to measure the CLV –such as expenditure rates, expenditure by visits to online and offline shops, word-of-mouth, and repeat purchases, among others–, few studies are focused on the optimization of the trade-off between the CLV and the marketing resources allocated by brands. Hence, our study aims at filling this gap.

4. Problem Formulation

Let us consider a set of potential customers I, and a set of potential commercial actions J. Regardless of the marketing plan, a customer will always be associated with one active action, the one with the highest impact on her CLV. Hence, assigning a customer $i \in I$ to an action $j \in J$ has an impact $b_{ij} \geq 0$ on the associated CLV. Also, deploying each action $j \in J$ has a cost $c_j > 0$, and there is a maximum budget $c_{max} \geq \min\{c_j\}$ available, i.e.: at least one marketing action can be deployed. Notice that, by varying the budget, it is possible to obtain a Pareto frontier of optimal CLV versus budget. In this context, the binary variable x_{ij} will take the value 1 if, and only if, customer i is assigned to action j, taking the the value 0 otherwise. Similarly, the binary variable y_j will take the value 1 if, and only if, action j is deployed, taking the value 0 otherwise. Now, the optimization problem can be formulated as follows:

$$\max \sum_{i \in I} \sum_{j \in J} b_{ij} x_{ij} \tag{1}$$

$$\text{s.t.:} \sum_{j \in J} c_j y_j \leq c_{max} \tag{2}$$

$$\sum_{j \in J} x_{ij} = 1 \quad \forall i \in I \tag{3}$$

$$x_{ij} \leq y_j \quad \forall i \in I, \quad \forall j \in J \tag{4}$$

$$x_{ij} \in \{0, 1\} \tag{5}$$

$$y_j \in \{0, 1\} \tag{6}$$

Equation (1) is the objective function, which consists in maximizing the aggregated CLV obtained by assigning each customer to one of the deployed commercial actions. Constraint (2) limits the total cost of the deployed actions to the available budget. Constraint (3) guarantees that each customer is assigned to exactly one commercial action. Constraint (4) ensures that a customer cannot be assigned to a non-active action. Finally, constraints (5) and (6) define the binary character of the decision variables.

5. A Biased-Randomized Algorithm

In order to solve the optimization problem described in Section 4, we first propose a constructive heuristic as described in Algorithm 1. The heuristic is designed to generate 'good' results –even for large instances– in extremely short computing times. It starts from a hypothetical scenario in which all possible actions proposed by the marketing department have been initially deployed. Notice that this scenario maximizes the CLV, but it might be unfeasible –since the cost of deploying all possible actions will typically exceed the available budget. Hence, it will be necessary to deactivate some of the proposed actions until we reach an affordable marketing cost. In order to select the right actions to deactivate, an efficiency criterion is needed. Thus, given a deployed action, $j \in J$, the *reduction in CLV* associated with its deactivation, $r_j > 0$, is computed as the difference between: *(i)* the current CLV of the customers assigned to j; and *(ii)* the CLV of the same set of customers once they have been reassigned to the best-fit alternative actions –notice that different customers might be reassigned to different active actions. It is possible then to define the *efficiency of deactivating* a deployed facility j, e_j, as a linear combination of two components. These are the inverse of the reduction in CLV, and the cost of deploying j, i.e.:

$$e_j = \frac{\alpha}{r_j} + (1 - \alpha) \cdot c_j, \tag{7}$$

where $\alpha \in (0, 1]$ is a parameter that depends on the specific inputs of the instance data and has to be determined experimentally. Notice that, in the particular case that all commercial actions have the same cost (i.e., $c_j = c, \forall j \in J$), then $\alpha = 1$ (in other words, only the inverse of the reduction in CLV will be considered in that particular case). We can now consider a list containing all deployed actions sorted from higher to lower efficiency. Then, the procedure will iteratively deactivate the next action in the list until the remaining actions deployed fit into our marketing budget.

The input parameters of the previously described heuristic algorithm are the following ones: a vector including the actions $j \in J$ and their associated deployment cost *(actions)*, a matrix containing the impact over the CLV associated with assigning each customer $i \in I$ to each commercial action $j \in J$ *(impacts)*, and the budget available to deploy commercial actions *(budget)*.

Algorithm 1 Constructive Heuristic

function `heuristic`(actions, impacts, budget)
sol $\leftarrow$ deployAllActions(actions, impacts) *% initial solution with max CLV*
for each action in actions **do**
    efficiency(action) $\leftarrow$ computeEfficiency(action, actions, impacts)
end for
effList $\leftarrow$ sortList(actions, efficiency)
while effList $\neq \emptyset$ **and** cost(sol) $>$ budget **do**
    nextAction $\leftarrow$ extract(effList, 0) *% extracts action in position 0*
    sol $\leftarrow$ deactivate(nextAction)
    for each action in effList **do**
        efficiency(action) $\leftarrow$ updateEfficiency(action, effList, impacts)
    end for
end while
return sol

As explained in Ferone et al. (2019), a constructive heuristic such as the previously described can be extended into a biased-randomized algorithm (Algorithm 2) by considering a multi-start framework and introducing the following updates: *(i)* an oriented random selection of the next action to be extracted from the efficiency list; and *(ii)* a local search process that helps to enhance the quality of the emerging solution. Regarding the first update, the idea is to assign decreasing probabilities of being selected to different elements in the sorted efficiency list. Thus, those actions with a higher efficiency will also be the ones receiving the higher probabilities, and vice versa. This strategy allows us to run the randomized heuristic multiple times, obtaining different solutions as a result –some of them of better quality than the original one provided by the constructive heuristic. For introducing this random behavior, we made use of a Geometric probability distribution with a parameter $\beta \in (0, 1)$. As explained in Belloso et al. (2019), values of β close to 0 lead to a uniform random selection of actions from the efficiency list. On the contrary, values of β close to 1 emulate the greedy behavior of the original heuristic. Finally, intermediate values explore different options by assigning higher probabilities of being selected to those actions showing higher efficiency levels. Regarding the local search, the idea is to improve the solution by exploring its neighborhood. In our case, we achieve this by randomly selecting two actions with different status (one deployed and one deactivated), and and then switching their current status –i.e.: deactivating the deployed one and vice versa. During each local search call, this swapping operation is iteratively performed for a specified number of times (a design parameter). Any time a better solution is found by applying the former local search, the current solution is updated.

Algorithm 2 Biased-Randomized Algorithm

```
function multi-start(actions, impacts, budget, maxTime, beta, nIter)
time ← 0
bestSol ← heuristic(actions, impacts, budget)
bestSol ← localSearch(bestSol, impacts, budget, nIter)
while time < maxTime do
    newSol ← heuristic(actions, impacts, budget, Geometric, beta) % randomized version of heuristic
    newSol ← localSearch(newSol, impacts, budget, nIter)
    if clv(newSol) > clv(bestSol) then
        bestSol ← newSol % choose the solution with the best customer lifetime value
    end if
    time ← updateTime
end while
return bestSol
```

6. An Illustrative Example

In order to facilitate a better understanding of the methodology presented in the previous section, we provide a numerical example, which is also supported by Figure 2. To illustrate this example, we have used a toy instance composed of 20 possible actions to be deployed, 20 customers, and a maximum budget of $\$90,623$ to invest in the candidate actions. As depicted in Figure 2a, our algorithm starts with an ideal scenario where all the possible actions are initially deployed, i.e, each customer receives her own fully customized action. Notice that, although this strategy maximizes the CLV, the associated cost exceeds the available budget. Thus, this initial solution is unfeasible. With the objective of generating a feasible solution, the constructive heuristic will create a list of actions sorted by efficiency and will iteratively deactivate actions until reaching a solution that satisfies the available budget. Figure 2b displays the first iteration of the algorithm, in which action 4 is deactivated, obtaining a CLV of 19.0 with an associated investment of $\$428,301$. Notice that each time an action is deactivated, the customers previously covered by it need to be re-assigned other actions (the ones that closely meet their interests among the deployed actions). In this case, customer 4 has been reassigned to action 2. Next, as the cost of the current solution still exceeds the available budget, the algorithm performs a new iteration. Figure 2c shows this second iteration, where action 13 has been deactivated and its associated customer has been reassigned to action 20. The resulting solution shows a CLV of 18.0 with a cost of $\$405,940$. This heuristic procedure will be repeated until a feasible solution is found. Figure 2d shows the final solution generated by the heuristic, which is a feasible one (its associated cost is $\$74,583$, which is lower than the budget). In this solution, just three actions are deployed, and all the customers have been assigned to one of these actions (the best one for them). However the CLV has decreased down to 3.07. In order to search for better solutions, the multi-start biased-randomized version of the algorithm is employed. Figure 2e shows a partial solution obtained by the biased-randomized algorithm after 5 seconds of computation. Finally, Figure 2f depicts the final solution after 10 seconds of computation. This last solution shows a total of 6 deployed actions, a CLV of 6.08, and a feasible associated cost of $\$83,428$.

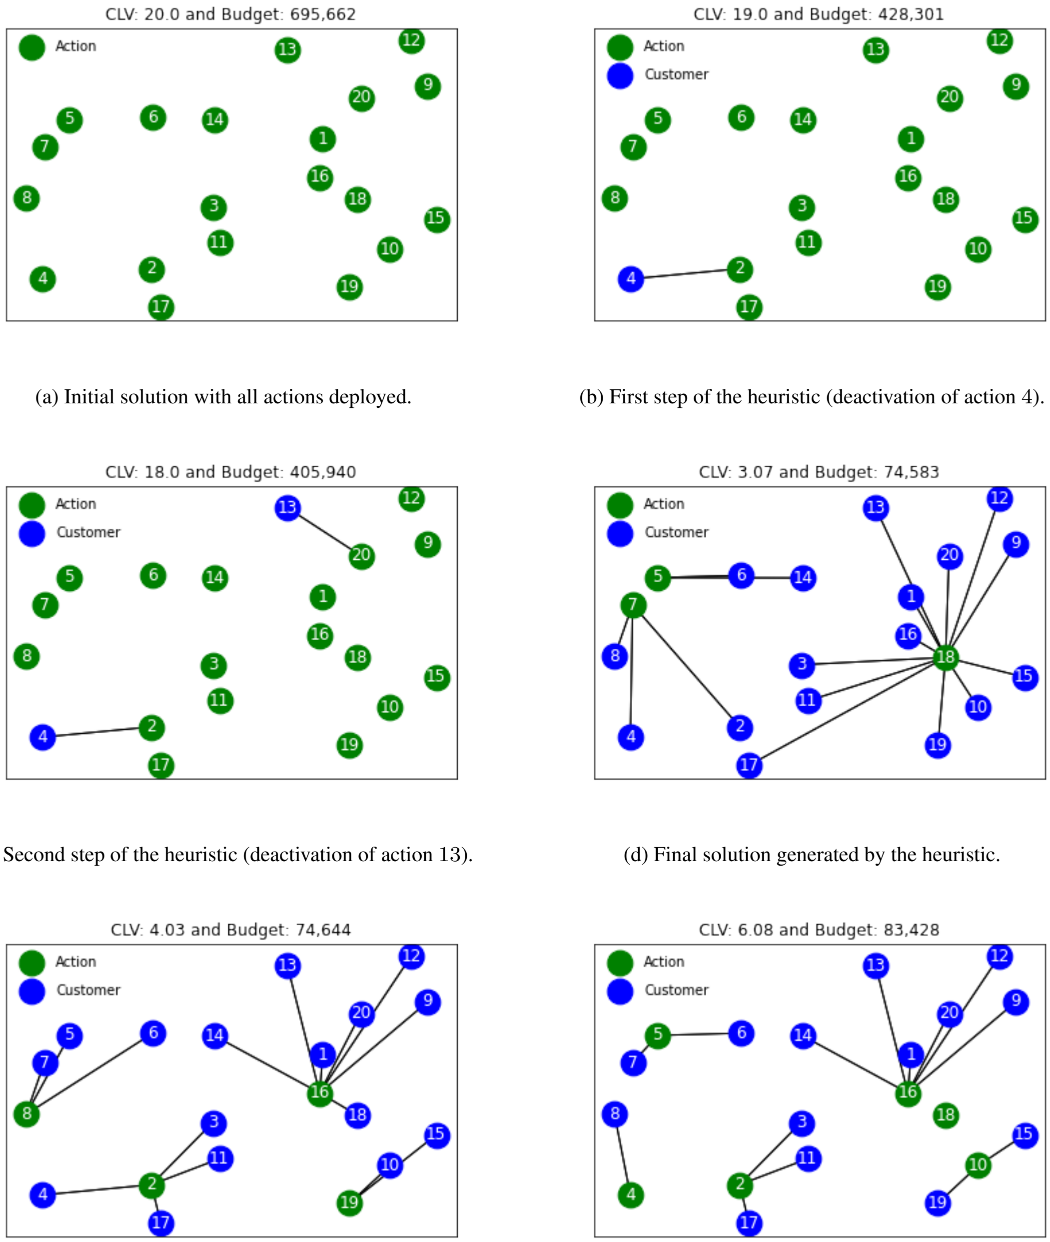

(a) Initial solution with all actions deployed.

(b) First step of the heuristic (deactivation of action 4).

(c) Second step of the heuristic (deactivation of action 13).

(d) Final solution generated by the heuristic.

(e) Solution by the biased-randomized algorithm after 5 seconds.

(f) Solution by the biased-randomized algorithm after 10 seconds.

Fig. 2: A case study illustrating the behavior of the proposed algorithm.

7. Computational results

The proposed approach has been implemented using Python 3.7 and tested on a workstation with a multi-core processor Intel Xeon E5-2650 v4 with 32GB of RAM. To the best of our knowledge, there are no benchmarks for the CLV problem. Accordingly, we have adapted the benchmark proposed by Ahn et al. (1988), which was originally designed for the *p*-median problem, and later used in the context of the facility location problem by Barahona and Chudak (2005). We have used the set of instances called MED, since they are the largest and most challenging ones. Each instance is composed of a set of *n* points, which are randomly chosen in the unit square. In our case, a point represents both a customer and a possible action, and the corresponding Euclidean distance determines how close a customer's needs / interests are from the corresponding action (i.e., as the distance between a customer and an action approaches to zero, the more personalized the latter is with respect to the former). The set consists of six different subsets, each with a different number of facilities and customers (500, 1000, 1500, 2000, 2500, and 3000), and three different opening cost schemes for each subset: $\sqrt{n}/10$, $\sqrt{n}/10$, $\sqrt{n}/100$, and $\sqrt{n}/1000$, corresponding to 10, 100, and 1000 instance suffixes, respectively. In order to adapt these instances to the CLV problem considered in this paper, we have carried out two additional modifications. Firstly, when a customer $i \in I$ is assigned to an action $j \in J$, the associated CLV is computed as follows:

$$b_{ij} = \frac{1}{1 + d_{ij}}, \tag{8}$$

where d_{ij} refers to the Euclidean distance between both points (notice that by increasing the distance in one unit, we guarantee that the denominator will never be zero). The main idea behind this equation is to consider that the impact of an action over a specific customer is inversely proportional to the distance between them, i.e.: a user will have a higher affinity for a campaign that is 'closer' to her needs and interests. As a second modification, we have included a new parameter to these instances. This parameter is the available budget, which has been set to 826,230 in our numerical experiments.

Table 1 reports the obtained results, where each row corresponds to a single instance. The first column identifies the instance. The next two columns show the optimal results, which were computed using the popular commercial solver IBM CPLEX. These values have been computed to explore the limits of the maximum problem size that can be solved using exact methods, and to validate the quality of the results provided by our approach. Notice that in the case of some instances (e.g., 2000-10), the optimal values cannot be computed, in a reasonable time and standard memory resources, by using exact methods. Hence, the importance of using heuristic-based approaches as the one proposed in this paper, specially when dealing with real-life problems or large size. The next two columns in the table show the obtained results provided by the basic constructive heuristic. Columns six and seven show the results provided when the local search operator is added to the constructive heuristic. Likewise, the next two columns show the results generated by the biased-randomized algorithm. Finally, the last three columns show the gap with respect to the optimal values, which are computed as indicated in Equation (9).

$$Gap\,(cost, cost_{optimal}) = 100 * \left(\frac{cost - cost_{optimal}}{cost_{optimal}}\right) \tag{9}$$

The results show that the biased-randomized approach provides high-quality competitive solutions for the CVL problem. Figure 3 depicts an overview of Table 1, which shows the performance of the different versions of our algorithms. In this box-plot, the vertical axis represents the gap obtained with respect to the optimum value, which represents the lower bound (LB). The results report that the greedy heuristic presents the highest gaps, since it is a simple and extremely fast method that is intended to support real-time decision making. On the average, it presents a gap of 7.9% with respect to the optimum values. Notice that it employs an average computational time of just 3.1 seconds for the Python implementation (which means that it can achieve an average far below 1 second when implemented in Java or C++). When the local search is added to the constructive heuristic, the average gap with respect to the optimal values is about 3.36%, which is noticeably lower than the previous one. Figure 4 illustrates the required computational times to obtain the best-found solutions. Notice that the local search is a fast operator, which increases the execution time with respect to the heuristic in just about 0.20 seconds. Hence, it seems clear that adding the local search to the heuristic is a good design option. Finally, when the heuristic is turned into a biased-randomized algorithm, our method is able to match –or remain very close to– the optimal solutions for all the instances, obtaining an average gap around 0.46%. These results highlight the effectiveness of our biased-randomized algorithm, which is able to provide near-optimal solutions in an average computational time of 81 seconds –notice that this is a very competitive time in comparison with the average time required by the exact method to reach the optimal values, which is above 1.5 hours.

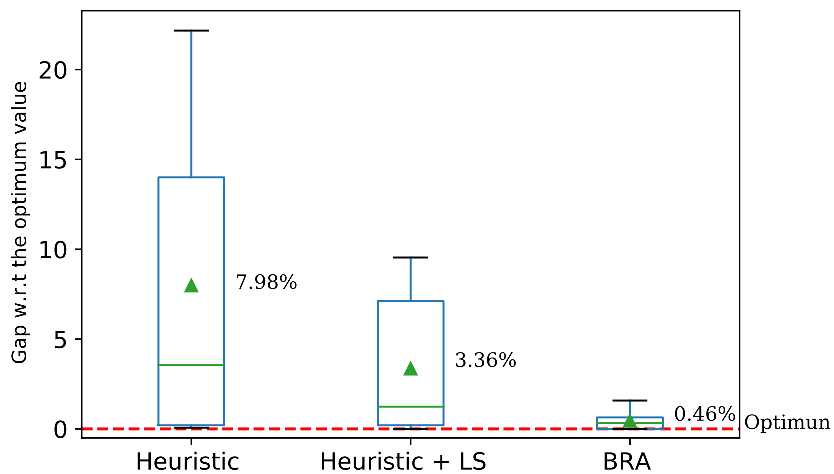

Fig. 3: A comparison of solutions quality for different approaches.

8. Conclusions

In the context of intelligent marketing, this paper has analyzed the concept of customer lifetime value. After reviewing existing work and definitions related to this concept, we have discussed several studies

Table 1: Results for the CLV - Comparing exact method (Cplex) with heuristic-based algorithm (Heuristic), heuristic combined with local search (Heuristic + LS) and the biased randomization algorithm (BRA).

	Cplex (Optimal)		Heuristic		Heuristic + LS		BRA		GAP(%)		
Instance	Reward [1]	Time(s)	Reward [2]	Time(s)	Reward [3]	Time(s)	Reward [4]	Time(s)	[1-2]	[1-3]	[1-4]
500-10	41.30	42.8	36.60	0.25	39.84	0.24	41.06	63.8	11.38	3.54	0.59
500-100	374.11	40.9	369.30	0.19	371.68	0.22	373.66	45.8	1.29	0.65	0.12
500-1000	500.00	21.7	499.00	0.17	499.02	0.22	500.00	0.9	0.20	0.20	0.00
1000-10	35.10	761.2	27.32	0.80	32.60	0.88	34.80	54.8	22.17	7.11	0.86
1000-100	273.02	347.2	263.32	0.78	269.63	0.89	271.27	125.1	3.55	1.24	0.64
1000-1000	1000.00	88.0	999.00	0.76	1000.00	0.80	1000.00	1.03	0.10	0.00	0.00
1500-10	44.27	790.6	35.83	1.81	40.04	3.02	44.12	40.58	19.06	9.54	0.32
1500-100	252.56	734.5	217.21	1.50	232.55	1.68	248.57	159.5	14.00	7.92	1.58
1500-1000	1500.00	242.7	1499.00	1.63	1500.00	1.68	1500.00	1.92	0.07	0.00	0.00
2000-10	-	-	36.41	4.61	36.41	4.14	38.65	52.42	-	-	-
2000-100	-	-	190.36	2.94	221.85	3.09	226.94	107.4	-	-	-
2000-1000	-	-	1,848.63	2.98	1858.23	2.98	1,861.90	152.5	-	-	-
2500-10	-	-	34.93	5.79	34.93	7.54	35.19	19.93	-	-	-
2500-100	-	-	178.72	5.02	230.13	4.77	235.35	152.46	-	-	-
2500-1000	-	-	1656.96	4.65	1674.39	5.00	1680.40	67.04	-	-	-
3000-10	-	-	32.28	7.83	33.40	10.68	35.52	51.11	-	-	-
3000-100	-	-	189.93	6.91	215.27	7.11	229.71	139.17	-	-	-
3000-1000	-	-	1515.69	8.10	1581.51	16.62	1590.60	21.51	-	-	-
Average:	446.71	341.1	535.03	3.15	548.42	3.98	552.65	69.8	7.98	3.36	0.46

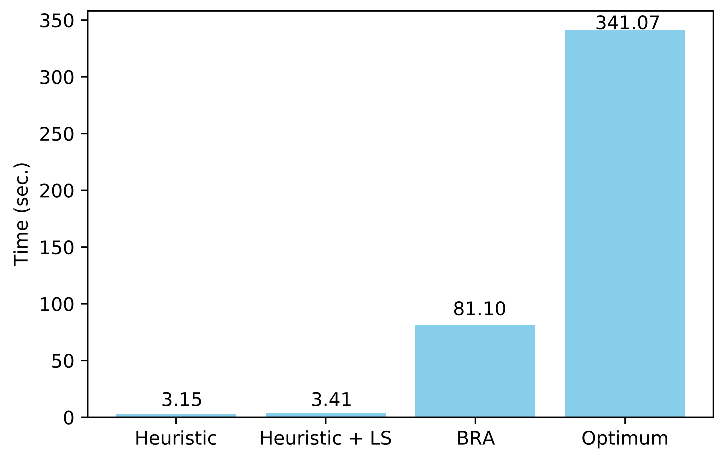

Fig. 4: A comparison of computational times for different approaches.

addressing the issue of maximizing this value under budget constraints. Then, we have proposed a model for this optimization problem. Our model is inspired by the uncapacitated facility location problem, with facilities playing the role of marketing actions, customers being represented by nodes, and the impact level of each action on each customer's CLV being defined by the inverse of the distance between the customer and the action (i.e., the closer an action to a customer, the more it can have a positive impact on her CLV).

Being a complex optimization problem, we propose a heuristic procedure to solve it in real time, even in the case of large-scale instances. This heuristic is later extended into a biased-randomized algorithm, which allows us to improve the quality of our results without significantly increasing computational times. The quality of the marketing strategies proposed to maximize lifetime values are validated by a direct comparison against optimal solutions generated after several hours of computation. As future work, we are considering the following possibilities: *(i)* to test our approach in a real-life scenario, contrasting our results with those obtained without using an algorithm-based approach; *(ii)* to consider random impact levels instead of deterministic ones to make the problem even more realistic; and *(iii)* to apply similar optimization approaches to other challenges in the fields of intelligent marketing and e-marketing.

Acknowledgments

This paper has been partially supported by the Catalan Department of Agriculture, Livestock, Fisheries and Food (56.21.003.2019.2A), AGAUR and FEDER.

References

AboElHamd, E., Shamma, H.M., Saleh, M., 2020. Dynamic programming models for maximizing customer lifetime value: An overview. In *Advances in Intelligent Systems and Computing*, Vol. 1037, Springer, pp. 419–445.

Ahn, S., Cooper, C., Cornuejols, G., Frieze, A., 1988. Probabilistic analysis of a relaxation for the k-median problem. *Mathematics of Operations Research* 13, 1, 1–31.

Barahona, F., Chudak, F.A., 2005. Near-optimal solutions to large-scale facility location problems. *Discrete Optimization* 2, 1, 35–50.

Belloso, J., Juan, A.A., Faulin, J., 2019. An iterative biased-randomized heuristic for the fleet size and mix vehicle-routing problem with backhauls. *International Transactions in Operational Research* 26, 1, 289–301.

Berger, P.D., Nasr, N.I., 1998. Customer lifetime value: Marketing models and applications. *Journal of Interactive Marketing* 12, 1, 17–30.

Blattberg, R.C., Malthouse, E.C., Neslin, S.A., 2009. Customer Lifetime Value: Empirical Generalizations and Some Conceptual Questions. *Journal of Interactive Marketing* 23, 2, 157–168.

Borle, S., Singh, S.S., Jain, D.C., 2008. Customer lifetime value measurement. *Management Science* 54, 1, 100–112.

Chang, W.J., 2020. Experiential marketing, brand image and brand loyalty: a case study of starbucks. *British Food Journal*

Ching, W.K., Ng, M.K., Wong, K.K., Altman, E., 2004. Customer lifetime value: stochastic optimization approach. *Journal of the Operational Research Society* 55, 8, 860–868.

Crowder, M., Hand, D.J., Krzanowski, W., 2007. On optimal intervention for customer lifetime value. *European Journal of Operational Research* 183, 3, 1550–1559.

Dominguez, O., Juan, A.A., Faulin, J., 2014. A biased-randomized algorithm for the two-dimensional vehicle routing problem with and without item rotations. *International Transactions in Operational Research* 21, 3, 375–398.

Dwyer, F.R., 1997. Customer lifetime valuation to support marketing decision making. *Journal of Interactive Marketing* 11, 4, 6–13.

Ehigie, B.O., 2006. Correlates of customer loyalty to their bank: A case study in Nigeria. *International Journal of Bank Marketing* 24, 7, 494–508.

Ekinci, Y., Ulengin, F., Uray, N., 2014a. Using customer lifetime value to plan optimal promotions. *Service Industries Journal* 34, 2, 103–122.

Ekinci, Y., Uray, N., Ülengin, F., 2014b. A customer lifetime value model for the banking industry: a guide to marketing actions. *European Journal of Marketing* 48, 3/4, 761–784.

Ferone, D., Gruler, A., Festa, P., Juan, A.A., 2019. Enhancing and extending the classical grasp framework with biased randomisation and simulation. *Journal of the Operational Research Society* 70, 8, 1362–1375.

Fruchter, G.E., Sigué, S.P., 2009. Social relationship and transactional marketing policies-maximizing customer lifetime value. *Journal of Optimization Theory and Applications* 142, 3, 469–492.

Ganiyu, R.A., Uche, I.I., Elizabeth, A.O., 2012. Is customer satisfaction an indicator of customer loyalty? *Australian Journal of Business and Management Research* 2, 7, 14–20.

Garepasha, A., Aali, S., Zendeh, A.R.B., Iranzadeh, S., 2020. Relationship dynamics in customer loyalty to online banking services. *Journal of Islamic Marketing*

Gupta, S., Hanssens, D., Hardie, B., Kahn, W., Kumar, V., Lin, N., Ravishanker, N., Sriram, S., 2006. Modeling customer lifetime value. *Journal of Service Research* 9, 2, 139–155.

Gupta, S., Lehmann, D.R., Stuart, J.A., 2004. Valuing customers. *Journal of marketing research* 41, 1, 7–18.

Haenlein, M., Kaplan, A.M., Schoder, D., 2006. Valuing the real option of abandoning unprofitable customers when calculating customer lifetime value. *Journal of Marketing* 70, 3, 5–20.

Hennig-Thurau, T., Langer, M.F., Hansen, U., 2001. Modeling and Managing Student Loyalty: An Approach Based on the Concept of Relationship Quality. *Journal of Service Research* 3, 4, 331–344.

Jonker, J.J., Piersma, N., Van Den Poel, D., 2004. Joint optimization of customer segmentation and marketing policy to maximize long-term profitability. *Expert Systems with Applications* 27, 2, 159–168.

Khajvand, M., Zolfaghar, K., Ashoori, S., Alizadeh, S., 2011. Estimating customer lifetime value based on RFM analysis of customer purchase behavior: Case study. *Procedia Computer Science* 3, 57–63.

Kumar, V., 2006. CLV: the databased approach. *Journal of Relationship Marketing* 5, 2-3, 7–35.

Kumar, V., Lemon, K.N., Parasuraman, A., 2006. Managing customers for value: An overview and research agenda. *Journal*

of Service Research 9, 2, 87–94.
Kumar, V., Ramani, G., Bohling, T., 2004. Customer lifetime value approaches and best practice applications. *Journal of Interactive Marketing* 18, 3, 60–72.
Kumar, V., Venkatesan, R., Bohling, T., Beckmann, D., 2008. The power of CLV: Managing customer lifetime value at IBM. *Marketing Science* 27, 4, 585–599.
Mokhtar, S.A., Katan, H., Imdadullah, H.u.R., 2018. Mobile banking adoption: the impacts of social influence, ubiquitous finance control and perceived trust on customers' loyalty. *Science International (Lahore)* 29, 4, 829–836.
von Mutius, B., Huchzermeier, A., 2021. Customized targeting strategies for category coupons to maximize clv and minimize cost. *Journal of Retailing*
Pagès-Bernaus, A., Ramalhinho, H., Juan, A.A., Calvet, L., 2019. Designing e-commerce supply chains: a stochastic facility–location approach. *International Transactions in Operational Research* 26, 2, 507–528.
Quintero-Araujo, C.L., Caballero-Villalobos, J.P., Juan, A.A., Montoya-Torres, J.R., 2017. A biased-randomized metaheuristic for the capacitated location routing problem. *International Transactions in Operational Research* 24, 5, 1079–1098.
Reichheld, F.F., Kenny, D.W., 1990. The hidden advantages of customer retention. *Journal of Retail Banking* 12, 4, 19–23.
Roemer, E., 2006. The impact of dependence on the assessment of customer lifetime value in buyer-seller relationships. *Journal of Marketing Management* 22, 1-2, 89–109.
Romero, J., van der Lans, R., Wierenga, B., 2013. A partially hidden markov model of customer dynamics for CLV measurement. *Journal of Interactive Marketing* 27, 3, 185–208.
Rust, R.T., Kumar, V., Venkatesan, R., 2011. Will the frog change into a prince? Predicting future customer profitability. *International Journal of Research in Marketing* 28, 4, 281–294.
Venkatesan, R., Kumar, V., 2004. A customer lifetime value framework for customer selection and resource allocation strategy. *Journal of Marketing* 68, 4, 106–125.
Verhoef, P.C., Lemon, K.N., 2013. Successful customer value management: Key lessons and emerging trends. *European Management Journal* 31, 1, 1–15.
Yi, Y., La, S., 2004. What Influences the Relationship between Customer Satisfaction and Repurchase Intention? Investigating the Effects of Adjusted Expectations and Customer Loyalty. *Psychology and Marketing* 21, 5, 351–373.
Zeithaml, V.A., Lemon, K.N., Rust, R.T., 2001. *Driving customer equity: How customer lifetime value is reshaping corporate strategy*. Simon and Schuster.