Constrained Multiobjective Optimization for the Design of Energy-Efficient Context Recognition Systems

Aljoša Vodopija[1,2](✉), Vito Janko[1,2], Mitja Luštrek[1,2], and Bogdan Filipič[1,2]

[1] Jožef Stefan Institute, Jamova cesta 39, 1000 Ljubljana, Slovenia
{aljosa.vodopija,vito.janko,mitja.lustrek,bogdan.filipic}@ijs.si
[2] Jožef Stefan International Postgraduate School, Jamova cesta 39, 1000 Ljubljana, Slovenia

Abstract. Context recognition (CR) systems infer the user's context, such as their physical activity, from sensor data obtained, for example, with smartphone sensors. Designing an energy-efficient CR system, however, is a complex optimization problem involving conflicting objectives and several constraints arising from real-world limitations and designers' preferences. To address this task, we propose a constrained multiobjective formulation of the CR design problem. Unlike most studies in this domain, we use a true multiobjective approach in solving it. Specifically, we apply a multiobjective evolutionary algorithm equipped with two different constraint handling techniques. Their performance is demonstrated in optimizing six CR systems of various complexity. The proposed problem formulation and the optimization results make it possible to better understand the CR systems operation and provide valuable information to the designers.

Keywords: Multiobjective optimization · Constraint handling · Context recognition · Energy efficiency

1 Introduction

Context recognition (CR) is a vague term encompassing a wide array of tasks where (usually wearable) sensors are used to detect something about the person wearing them. Possible applications range from counting steps, localization, detecting activities such as walking or running, to monitoring someone's physical and mental health.

CR is an already mature research area [13] and many applications using CR systems come pre-installed on average smartphones. However, a common problem that occurs when designing such systems is the energy consumption of the device that is collecting and processing the sensor data. It is easy to imagine that a smartphone application that uses all its sensors (e.g., GPS, Bluetooth,

B. Filipič et al. (Eds.): BIOMA 2020, LNCS 12438, pp. 308–320, 2020.
https://doi.org/10.1007/978-3-030-63710-1_24

Wi-Fi, accelerometer, etc.) can detect much about its user, but also quickly drains the phone's battery, making it useless in practice.

There are many ways of preserving the battery life of a CR system. One of the most effective ways is the choice of the right sensors for the task (as different sensors can be used for the same CR task) and duty-cycling them, e.g., periodically turning them on and off again. Energy savings can be further increased if the sensors used and the duty-cycle durations adapt to the current context. For example, one might want to use GPS when the user is driving, but accelerometer when walking.

The issue with creating such adaptive CR systems is that doing so requires either a lot of expert knowledge of the domain or manual experimentation. Thus, any process that could at least partially automate the task of searching for energy-efficient solutions would be greatly beneficial.

Janko et al. [8] were the first to show that this problem can be formulated as a multiobjective optimization problem (MOP) with the objectives being the accuracy and energy consumption of the CR system. Their work, however, lacked a thorough experimentation in solving the resulting MOP and did not consider constraints in its formulation. The constraints naturally arise from real-world limitations of some sensors and from additional desires from the system designers.

In this work we expand on both of these aspects by performing a more comprehensive experimental evaluation, and more importantly, adding real-world constraints to the proposed MOP. The resulting constrained MOP is solved using the well-known Nondominated Sorting Genetic Algorithm II (NSGA-II) [2]. Two constraint handling techniques (CHTs) are applied: the original constrained-domination principle (CDP) [2] and a more recent approach based on an ensemble (ENS) of multiple CHTs proposed in our previous work [12]. Their performance in solving the CR optimization problem is assessed on six progressively harder CR systems.

We first present two different datasets—Commodity12 and Opportunity—that represent two different CR problems (Sect. 2). In Sect. 3, we then elaborate on how to represent the semantics of these datasets as a MOP. Special consideration is given to the constraint formulation (Sect. 3.1), and for each dataset we prepare three different, progressively harder, sets of these constraints. In Sect. 4, we test the difficulty of the proposed CR optimization problems and evaluate the quality of the found energy-efficient solutions, and finally conclude in Sect. 5.

2 Datasets

In this section, we present two datasets from two CR problems. They both contain streams of sensor data, which are then split into windows and can be used to calculate features. These features are then fed into machine-learning classifiers whose goal is to classify each window into one of the predetermined contexts as accurately as possible and with as little sensor data as possible.

2.1 Commodity12

The aim of the Commodity12 project was to create a system that can be used by diabetics to monitor their activities and help them manage their lifestyle more easily. All details can be found in the previous work on the domain [1].

For data collection, a smartphone and a chest-worn heart-rate monitor were used to monitor ten participants. Each participant continuously collected data for two weeks and manually labeled the following contexts: *sleep*, *work*, *home*, *eating*, *transport*, *exercise*, *out* (out of house, but not in any of the previous contexts). The data was collected from ten sensors: accelerometer, barometer, light sensor, GPS location, a list of visible Wi-Fi networks, a description of location by the Foursquare web service, sound, time, heart rate and respiration rate. The first eight were measured with the smartphone, and the last two with the heart rate monitor connected to the smartphone via Bluetooth.

Random Forest was identified as the best-performing classifier and was therefore selected for the present work on this dataset. While the classification accuracy was reasonably high (between 73% and 88%, depending on the user), the energy consumption made the application impractical to use—and thus the need for energy optimization.

To use energy consumption as one of the optimization objectives, it needs to be estimated for each sensor combination (as the energy consumption of different sensors do not add up linearly). This was done empirically by attaching a multimeter device directly to the smartphone battery [8].

2.2 Opportunity

Opportunity [9] is a popular publicly available dataset designed to evaluate algorithms for detecting human activity. Data on four users were recorded while they were performing various tasks in an apartment.

There were 30 sensor clusters in this apartment, some on the user's body and some on the objects the user interacted with. The complete list of sensor locations is as follows: user's left knee, left and right upper arm, left and right forearm, user's hips, left and right shoe, left and right wrist, left and right hand, as well as a cup, salami, water bottle, cheese, bread, knife, sugar, plate, and drinking glass. Each cluster contained some of the following sensors: accelerometer, gyroscope and magnetometer.

The dataset provides various sets of labels, out of which we decided to test the case where the problem was to recognize which object the user is currently holding in their right hand. There were 18 classes: *bottle*, *bread*, *chair*, *cheese*, *cup*, *dishwasher*, *door*, *drawer*, *fridge*, *glass*, *knife*, *milk*, *plate*, *salami*, *spoon*, *sugar*, *switch*, *table*, *none*), each representing an object held, except for the *none* class that represented no object in hand. The class distribution was highly unbalanced with the *none* class having a representation of 57%.

The classification process was made relatively simple in order to conform to the introductory paper [10] of the dataset. The data from each sensor was divided into 500-ms non-overlapping windows on which we calculated the mean and standard deviation. The k-nearest neighbors classifier ($k = 3$) was then used for the classification.

This problem domain has an unusually high number of sensors (30 sensor clusters), which creates an enormous space for possible sensor subsets. Therefore we decided to use only some sensor subsets (not all, as with Commodity12) as the search-space for multiobjective optimization. These subsets were selected in the following way. We started with an empty set. Then we added the sensor cluster that increased the F-score the most to the set. This was repeated until no single sensor cluster could increase the F-score. Each resulting subset was added as a sensor setting (each subset had one more sensor than the previous one). The procedure was then repeated for each context, this time adding sensors only if it increased the F-score for recognizing this context. All generated subsets were added again as sensor settings. The justification for this greedy procedure is that most sensor subsets are redundantly large, both inflating the energy consumption and unnecessarily increasing the search space of different system configurations.

For the sake of simplicity (and since we did not have access to the details about the sensors) we assumed that all existing sensors had similar energy consumption. To model their combined consumption, we simply added up the individual energy consumptions.

3 Problem Formulation

Suppose the CR system can detect c different contexts. It can do so by using different settings—the setting being which sensors to use and with which duty-cycle schedules (sensors can work for a time periods, then sleep for s time periods, and repeat). Whenever a context is detected, the setting used changes to the one assigned for the current context (e.g., whenever *transport* is detected, the GPS gets turned on). This opens up the problem of finding the ideal assignment of each context to the one of the possible settings. Each such assignment will result in a different CR system that will generally have a different trade-off between its accuracy and its energy consumption. We can assume that both of these objectives can be accurately estimated using either a simulation or a mathematical model [6–8].

The problem can be naturally formulated as a multiobjective optimization problem with the accuracy of the system, f_1, and its energy consumption, f_2, being two conflicting objectives. A setting-to-context assignment can be represented with an integer (*decision*) *vector*,

$$x = (x_1, \ldots, x_D)^{\mathrm{T}} \in S \subset \mathbb{N}^D$$

where S denotes the *decision space* of dimension $D = 2c+1$. The first c entries of x dictate which sensor subset to use when the corresponding context is detected (possible sensors subsets are enumerated). Similarly, the second c entries dictate for how long the system sleeps in each duty cycle (no sensor is working). Finally, the last component indicates how long the sensors are active between the sleeping periods. It is of note that the length of a duty cycle is not fixed, therefore the lengths of the sleeping and active periods do not necessarily sum into a given total. Two duty cycles of different lengths may have different performance, even with the same ratio of active and sleeping periods.

The number of possible sensor subsets was roughly 200 for both datasets, while the lengths of both the sleeping and active periods were capped at 30. The ranges of these parameters were chosen to be semantically sensible and, in the case of no constraints, all parameter values have the potential to be part of a Pareto-optimal solution. The fitness of these integer vectors was calculated using the mathematical model from [7,8].

To make it possible to compare the performance of various CR systems, we consider normalized objective values. The values of f_1 are already normalized since they represent the achieved accuracy. On the other hand, the values of f_2 are normalized by the maximum possible energy consumption. This is obtained when all the sensors are used and they are never turned off.

3.1 Constraints

For both the Commodity12 and Opportunity datasets we derived three versions of constraints, each progressively harder than the previous one. The difficulty was increased either by adding additional constraints or by making the existing ones harder to satisfy. In the latter case we used the variable z to denote the value that was changing from one problem version to another. The used values of z for each problem setting are summarized in Table 1.

The first category of constraints is based on the precisions and recalls of specific contexts when the system is using a particular solution. In each dataset, we selected a subset of contexts (denoted as L) that represents contexts important for the real-life application of the system. In the Commodity12 problem the system has to give diabetic patients recommendations about their lifestyle, thus the most important contexts are: *eating*, *exercise* and *transport* (as it includes walking). In the Opportunity problem we wanted to detect the preparation of a sandwich, so the crucial contexts are: *bread*, *salami*, *plate*, *knife*, *fridge*, *drawer* and *none*.

For each of these contexts we wanted to ensure that their precision and recall do not significantly deviate from their maximum possible values (M_i). The maximum values are achieved when all the sensors are used and are never turned off (duty-cycled).

$$g_{1,i}(x) = \text{ precision}(i, x) \geq z \cdot M_i, \;\; i \in L \tag{1}$$

$$g_{2,i}(x) = \text{ recall}(i, x) \geq z \cdot M_i, \;\; i \in L \tag{2}$$

Here, recall(x, i) is the recall of the i-th activity when the system is using solution x, and z is a fraction that varies from problem to problem.

For other contexts, we still wanted that they are "balanced" and that the system is not entirely omitting one in favor of the others. Thus, the next set of constraints ensures that the precisions and recalls of these contexts are in a certain range from each other.

$$g_3(x) = \max\{\text{precision}(i, x) \mid i \notin L\} - \min\{\text{precision}(i, x) \mid i \notin L\} \leq 0.25 \quad (3)$$

$$g_4(x) = \max\{\text{recall}(i, x) \mid i \notin L\} - \min\{\text{recall}(i, x) \mid i \notin L\} \leq 0.25 \quad (4)$$

In many domains it has been shown [5] that the accuracy of the system can be improved by "smoothing" the predictions, i.e., classifying a few consecutive data windows and then taking the most frequent prediction for that time period. To allow for this post-processing step, we try to enforce a longer active period if the accuracy of the system is below some threshold.

$$g_5(x) = \begin{cases} x_{2c+1} \geq 5, & 0.5 < f_1(x) < 0.75 \\ x_{2c+1} \geq 3, & f_1(x) \leq 0.5 \\ x_{2c+1} \geq 1, & \text{otherwise} \end{cases} \quad (5)$$

Our duty-cycle scheme assumes that sensors can be switched on and off in short intervals, and can do so without any additional energy cost. This is frequently not the case and it creates additional constraints on the system design. For example, if the GPS is active, the sleeping part of the duty cycle has to be longer to account for the extra time needed for turning the GPS on and off again. In Eq. (6) used for the Commodity12 problem, we used binary variables, x_i^s, that indicate if sensor s is active when using the sensor set x_i (g stands for GPS, b for sensors that use Bluetooth and w for Wi-Fi). For the Opportunity problem we used a similar scheme, but made different weights based on whether the sensor is on the body or in the environment.

$$g_{6,i}(x) = \begin{cases} x_{i+c} \geq 8 + z, & x_i^g \\ x_{i+c} \geq 5 + z, & \neg x_i^g \wedge x_i^b \\ x_{i+c} \geq 3 + z, & \neg x_i^g \wedge \neg x_i^b \wedge x_i^w \\ x_{i+c} \geq 0, & \text{otherwise} \end{cases} \qquad i \in \{1, \ldots, c\} \quad (6)$$

The final constraint arises from the number of sensors being used by the system, as ideally we would like to use as few sensors as possible. Doing so in the case of Opportunity would mean reducing the cost of the hardware, while in Commodity12 it would reduce the number of different data types that system designers have to analyze. In the Opportunity problem we also want to limit the number of sensors worn by the user to increase the practicality of the system.

$$g_7(x) = |\bigcup_{i=1}^{c} \text{sens}(x_i)| \leq z \quad (7)$$

$$g_8(x) = |\bigcup_{i=1}^{c} \text{bsens}(x_i)| \leq z \quad (8)$$

Table 1. Values of z for each problem/constraint combination and the characteristics of the resulting test CR systems: the number of contexts c, dimension of the decision space D, and number of constraints N. If the parameter z is not used, the sign $+/-$ denotes whether the given constraint category is used ($+$) or not ($-$). In the case of OPP3 and constraints $g_{1,i}$ and $g_{2,i}$, all contexts have bounded precision and recall, not only the crucial ones.

System	$g_{1,i}$	$g_{2,i}$	g_3	g_4	g_5	$g_{6,i}$	g_7	g_8	c	D	N
COM1	0.8	0.8	$-$	$+$	$+$	-2	$-$	$-$	7	15	15
COM2	0.9	0.9	$-$	$+$	$+$	0	$-$	$-$	7	15	15
COM3	0.8	0.9	$+$	$+$	$+$	0	5	$-$	7	15	17
OPP1	0.7	0.7	$-$	$-$	$+$	-1	18	10	18	37	35
OPP2	0.8	0.8	$-$	$-$	$+$	-1	18	10	18	37	35
OPP3	0.9	0.9	$-$	$-$	$+$	0	18	10	18	37	57

Here, $\text{sens}(x_i)$ is the set of all sensors used by x_i, and $\text{bsens}(x_i)$ the set of all body-worn sensors used.

Throughout the paper we use COM as the abbreviation for Commodity12 test CR systems and OPP for Opportunity test CR systems. The characteristics of the test CR systems are summarized in Table 1. Additionally, we provide the feasibility ratio (the proportion of feasible solutions) of each optimization problem. The estimation is based on two samples of 10^6 solutions generated by random sampling and Latin hypercube sampling. The feasibility ratio for COM1 is approximately $5.3 \cdot 10^{-5}$ according to random sampling and $6.1 \cdot 10^{-5}$ according to Latin hypercube sampling. On the other hand, no feasible solutions can be found for other test CR systems regardless of the sampling method used. Therefore, their feasibility ratios are estimated to be less than 10^{-6}. Particularly hard constraints are g_3, g_7, and g_8 that are each satisfied in less than 1% of the sampled solutions.

4 Experiments and Results

Based on the multiobjective formulation of the CR optimization problem, the experimental evaluation aimed at finding sets of trade-off solutions in the form of Pareto front approximations. For this purpose we used the well-known NSGA-II multiobjective optimization algorithm equipped with CDP [2] and ENS [12].

The CDP technique is the most frequently used method to solve constrained MOPs in practice. It strictly favors feasible solutions over infeasible ones. While feasible solutions are ranked based on Pareto dominance, the infeasible solutions are ranked according to constraint violations.

The ENS method combines multiple CHTs into an ensemble-based method where solutions for a new generation are selected based on a weighted voting provided by various CHTs. This approach considers only CHTs which are applied in the replacement phase, i.e., survivor selection, of an evolutionary algorithm.

Table 2. Average cumulative hypervolume values obtained by both CHTs on the test CR systems.

System	CDP $[\mu \pm \sigma]$	ENS $[\mu \pm \sigma]$
COM1	1.0081 ± 0.0018	1.0215 ± 0.0028
COM2	0.9497 ± 0.0095	0.9517 ± 0.0111
COM3	0.8873 ± 0.0277	0.8905 ± 0.0269
OPP1	0.8011 ± 0.0054	0.8419 ± 0.0047
OPP2	0.7111 ± 0.0189	0.7614 ± 0.0127
OPP3	0.6131 ± 0.0199	0.6705 ± 0.0141

Each CHT in the ensemble is supposed to provide a quality measure combining individuals' objective values and constraint violations. These quality measures are normalized to allow for comparison of individuals' quality among various CHTs. The quality measure produced by the ensemble of CHTs is a weighted average of the corresponding quality measures.

In this work, four CHTs were considered for the ensemble: normalized overall constraint violation [11], CDP, dynamic penalty function [3], and multiple constraint ranking [4]. In contrast to the original work [12], we decided to change the nondominated sorting with the normalized overall constraint violation, since the proposed test CR systems are heavily constrained.

The experimental setup was defined in the following way. Both methods were run with populations of 200 solutions for 1000 generations. The crossover probability was set to 0.9 and the mutation probability to 0.1. These parameter values were selected based on the experimental results from [6,8]. Specifically, for ENS, uniform weights ($w_i = 1/4$ for $i \in \{1, 2, 3, 4\}$) were used, while the two parameters of the dynamic penalty function, C and α, were set to 0.5 and 2, respectively. On each test CR system, every CHT was run 31 times, each time with a new randomly initialized population.

Additionally, the implementation details and parameter settings concerning data preprocessing, feature extraction, Random Forest classifier learning, and calculation of energy consumption were defined as in [6].

The quality of the optimization algorithm runs was measured with the cumulative hypervolume of the Pareto front approximation found in each run. Given $f_1, f_2 \in [0, 1]$, the reference point for hypervolume calculations was set to $(-0.1, 1.1)^{\mathrm{T}}$.

The means of cumulative hypervolume values are shown in Table 2. As we can see, ENS obtains better cumulative hypervolume means than CDP on all test CR systems. However, the differences are negligible on both COM2 and COM3. Indeed, the independent Welch's t-test (the normality assumption was confirmed by the Shapiro-Wilk test, while the homoscedasticity was rejected by the Levene's test) shows statistically significant differences in algorithm performance for COM1, OPP1, OPP2 and OPP3 ($p < 0.05$), while there are no significant differences observed on COM2 and COM3 ($p \approx 0.24, 0.47$).

The results are even easier to interpret through visualization of the obtained Pareto front approximations. Figure 1 shows Pareto front approximations for the test CR systems resulting from typical runs. In more detail, all the runs corresponding to a given test CR system are sorted based on the obtained cumulative hypervolume, and the front obtained in the median run is shown in the figure. We can see that the fronts obtained by ENS are superior in both convergence and diversity. This is especially true on OPP test CR systems, where ENS obtains significantly better Pareto-optimal solutions than CDP. It is worth noting that the performance of CDP compared to ENS decreases with constraint complexity.

Interestingly, on COM2, a few solutions obtained by CDP dominate the solutions obtained by ENS (see Fig. 1, COM2, around $f_2 \approx 0.4$) although its front seems to be well converged. This observation suggests that ENS gets stuck in a sub-optimal region and reveals the problem's multimodal nature. Nevertheless, further investigation is needed to explain this phenomenon. Another interesting observation is the sharp knee appearing in the fronts for all COM test CR systems. Investigating the found solutions revealed that solutions on one side of the knee only use sensors for one time period in every duty cycle (and thus have low energy consumption), while the solutions on the other side have an increasingly longer active period. Finally, in all cases the energy consumption quickly drops (in exchange for a small accuracy loss), indicating that smaller sensor subsets can be almost as effective as all sensors.

Figure 2 shows the progress of the mean cumulative hypervolume during optimization for the test CR systems. The x-axis indicates the spent function evaluations and y-axis the corresponding cumulative hypervolume values. Although the performance of CDP and ENS are comparable on COM test CR systems, we can see that ENS is more efficient. On average ENS needs less function evaluations to converge than CDP, and this gap increases for more constrained CR systems. In addition, the graphs show that both CHTs converge on all test CR systems except on OPP3. For this reason, it is unlikely that an increase in the computational budget would drastically improve these results (except on OPP3).

Finally, since CR optimization is a design problem, the results with respect to efficiency (spent computational resources) are not of great importance. The most computationally expensive task is solution evaluation. A single solution evaluation takes around 0.016 s for COM test CR systems, and 0.217 s for OPP test CR systems. All the experiments were run on a 3.40 GHz Intel(R) Core(TM) i7-6700 CPU with 16 GB RAM.

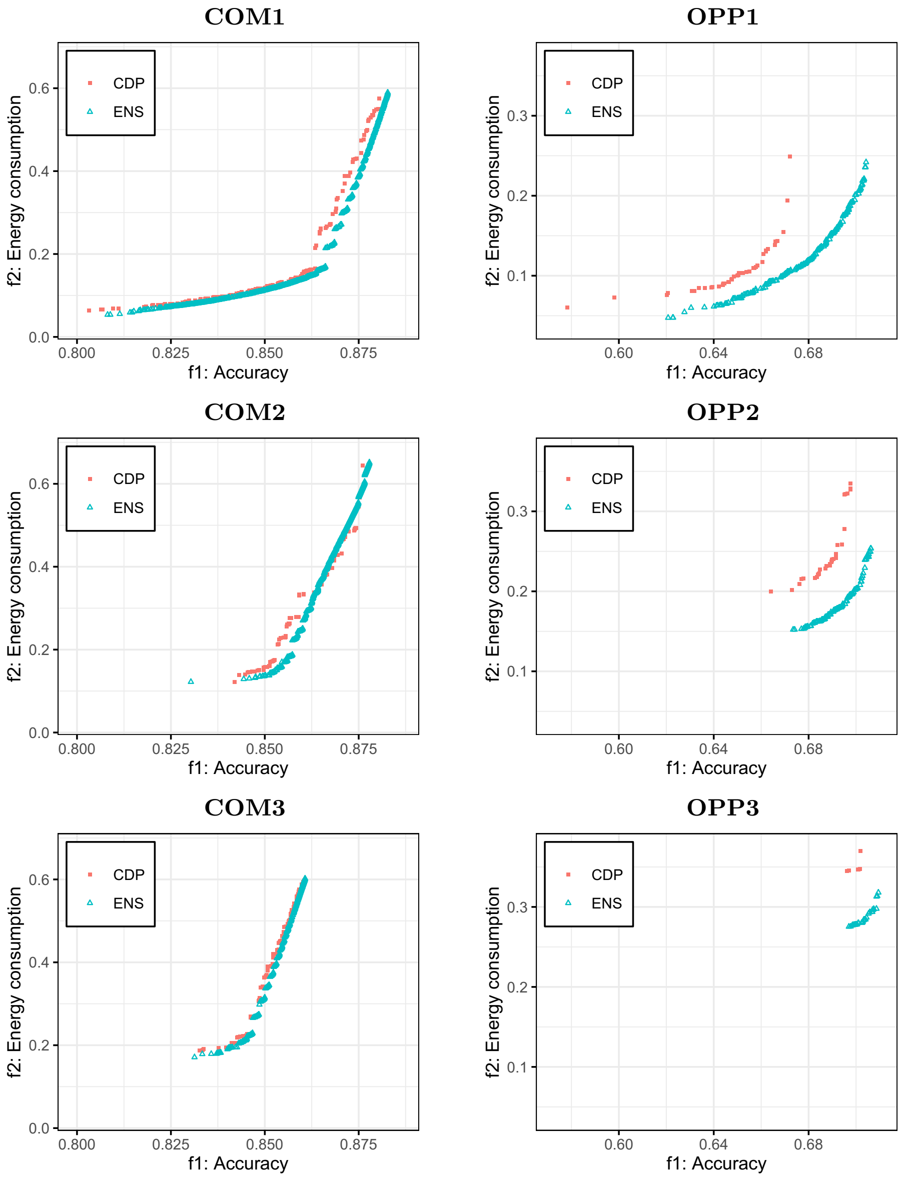

Fig. 1. Pareto front approximations for COM (left) and OPP (right) test CR systems.

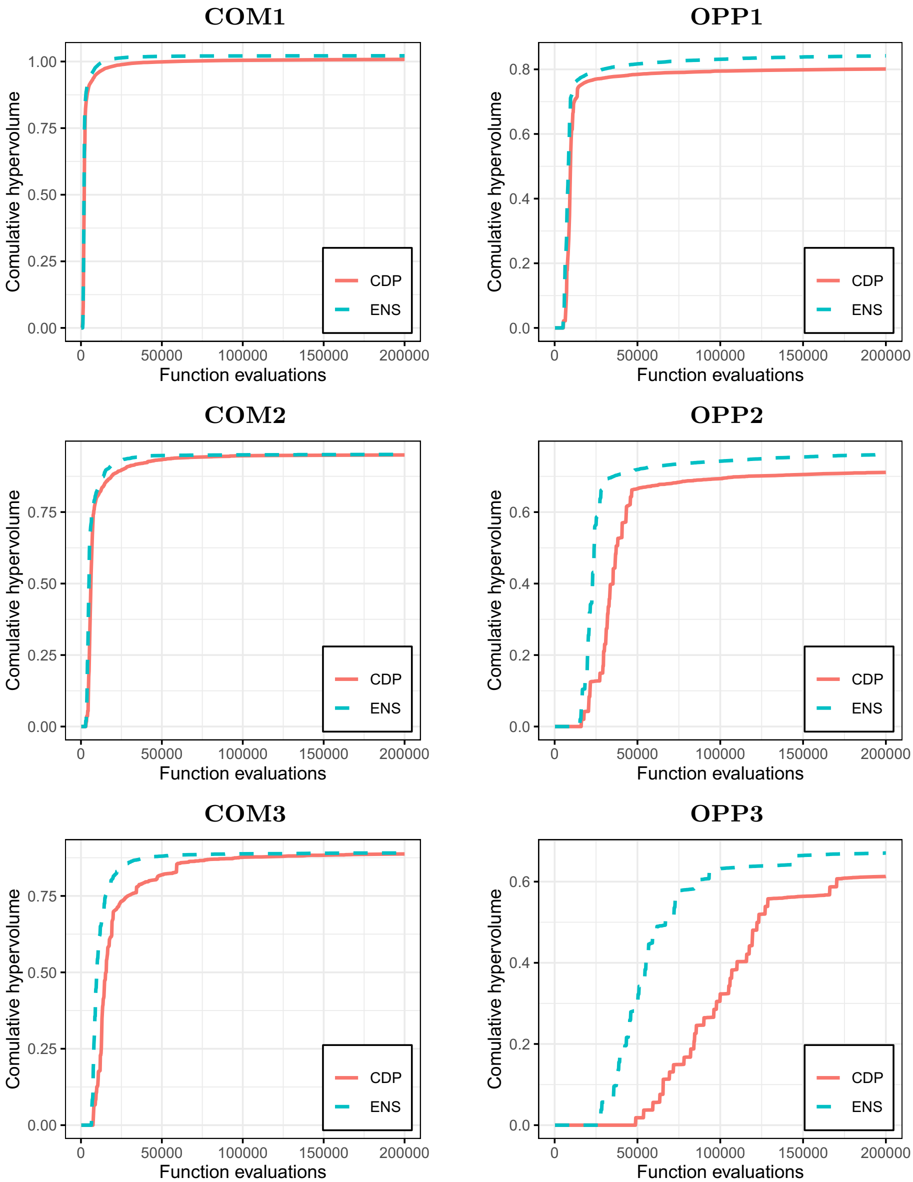

Fig. 2. Cumulative hypervolume progress for COM (left) and OPP (right) test CR systems.

5 Conclusions

In this paper, we expanded the work of Janko et al. [6] by proposing a constrained multiobjective optimization problem formulation for the design of energy-efficient CR systems. The proposed CR optimization problem takes into account the accuracy and overall energy consumption of the CR system and, at the same time, considers real-world limitations and designers' preferences. As opposed to most related work, the resulting optimization problem was solved using a true multiobjective optimizer capable of finding approximations of Pareto-optimal solutions. Specifically, the constraints were handled both by a classic technique frequently used in constrained multiobjective optimization, and our novel ensemble-based approach.

The experimental results on six progressively harder test CR systems show that the approach based on the ensemble paradigm performs better than the classic technique. The ensemble was superior on four test CR systems, while no differences in performance were observed on two easier CR systems. Additionally, an initial investigation of the produced Pareto front approximations reveals the multimodal nature of the CR optimization problem.

The found solutions were semantically meaningful as well as energy-efficient, especially in comparison to the base case where all the sensors were used. As an example, the "knee" solution for the COM1 test system represents a trade-off where, by sacrificing less than 2% of classification accuracy, the energy consumption is reduced by 82%.

In the future, we plan to investigate the CR optimization problem in more detail and assess the scalability of the applied optimization methodology. For the first task, we will examine the landscapes of the introduced optimization problem by investigating the produced solutions. For the second task, we will design new test CR systems, preferably using new datasets. Finally, the test CR systems will be made publicly available to the optimization community.

Acknowledgment. This work is part of a project that has received funding from the *European Union's Horizon 2020 research and innovation program* under Grant Agreement no. 692286. We also acknowledge financial support from the Slovenian Research Agency (young researcher program and research core funding no. P2-0209).

References

1. Cvetković, B., Janko, V., Romero, A.E., Kafalı, Ö., Stathis, K., Luštrek, M.: Activity recognition for diabetic patients using a smartphone. J. Med. Syst. **40**(12), 1–8 (2016). https://doi.org/10.1007/s10916-016-0598-y
2. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans. Evol. Comput. **6**(2), 182–197 (2002). https://doi.org/10.1109/4235.996017
3. Eiben, A.E., Smith, J.E.: Introduction to Evolutionary Computing. NCS. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-44874-8

4. Garcia, R.d.P., de Lima, B.S.L.P., Lemonge, A.C.d.C., Jacob, B.P.: A rank-based constraint handling technique for engineering design optimization problems solved by genetic algorithms. Comput. Struct. **187**, 77–87 (2017). https://doi.org/10.1016/j.compstruc.2017.03.023
5. Gjoreski, M., et al.: Classical and deep learning methods for recognizing human activities and modes of transportation with smartphone sensors. Inf. Fusion **62**, 47–62 (2020). https://doi.org/10.1016/j.inffus.2020.04.004
6. Janko, V.: Adapting sensor settings for energy-efficient context recognition. Ph.D. thesis, Jožef Stefan International Postgraduate School (2020). https://www.researchgate.net/publication/343230379
7. Janko, V., Luštrek, M.: Choosing duty-cycle parameters for context recognition. In: 14th International Conference on Intelligent Environments, pp. 83–86. IEEE (2018). https://doi.org/10.1109/IE.2018.00020
8. Janko, V., Luštrek, M.: Using Markov chains and multi-objective optimization for energy-efficient context recognition. Sensors **18**(1), 80 (2018). https://doi.org/10.3390/s18010080
9. Roggen, D., et al.: Collecting complex activity datasets in highly rich networked sensor environments. In: Seventh International Conference on Networked Sensing Systems (INSS), pp. 233–240. IEEE (2010). https://doi.org/10.1109/INSS.2010.5573462
10. Sagha, H., et al.: Benchmarking classification techniques using the Opportunity human activity dataset. In: IEEE International Conference on Systems, Man, and Cybernetics, pp. 36–40. IEEE (2011). https://doi.org/10.1109/ICSMC.2011.6083628
11. Tanabe, R., Oyama, A.: A note on constrained multi-objective optimization benchmark problems. In: IEEE Congress on Evolutionary Computation (CEC), pp. 1127–1134. IEEE (2017). https://doi.org/10.1109/CEC.2017.7969433
12. Vodopija, A., Oyama, A., Filipič, B.: Ensemble-based constraint handling in multiobjective optimization. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO 2019), pp. 2072–2075. ACM (2019). https://doi.org/10.1145/3319619.3326909
13. Yürür, Ö., Liu, C.H., Sheng, Z., Leung, V.C., Moreno, W., Leung, K.K.: Context-awareness for mobile sensing: a survey and future directions. IEEE Commun. Surv. Tutor. **18**(1), 68–93 (2016). https://doi.org/10.1109/COMST.2014.2381246