Largest Small *n*-Polygons: Numerical Results and Conjectured Optima

János D. Pintér
Department of Industrial and Systems Engineering
Lehigh University, Bethlehem, PA, USA
jdp416@lehigh.edu

Abstract

LSP(*n*), the *largest small polygon* with *n* vertices, is defined as the polygon of unit diameter that has maximal area *A*(*n*). Finding the configuration *LSP*(*n*) and the corresponding *A*(*n*) for even values $n \geq 6$ is a long-standing challenge that leads to an interesting class of nonlinear optimization problems. We present numerical solution estimates for all even values $6 \leq n \leq 80$, using the AMPL model development environment with the LGO nonlinear solver engine option. Our results compare favorably to the results obtained by other researchers who solved the problem using exact approaches (for $6 \leq n \leq 16$), or general purpose numerical optimization software (for selected values from the range $6 \leq n \leq 100$) using various local nonlinear solvers. Based on the results obtained, we also provide a regression model based estimate of the optimal area sequence $\{A(n)\}$ for $n \geq 6$.

Key words

Largest Small Polygons · Mathematical Model · Analytical and Numerical Solution Approaches · AMPL Modeling Environment · LGO Solver Suite For Nonlinear Optimization · AMPL-LGO Numerical Results · Comparison to Earlier Results · Regression Model Based Optimum Estimates

1 Introduction

The diameter of a (convex planar) polygon is defined as the maximal distance among the distances measured between all vertex pairs. In other words, the diameter of the polygon is the length of its longest diagonal. The *largest small polygon* with *n* vertices is the polygon of unit diameter that has maximal area. For any given integer $n \geq 3$, we will refer to this polygon as *LSP*(*n*) with area *A*(*n*). For illustration, see Figure 1 – cited from the referenced webpages of Weisstein – that shows the largest small hexagon *LSP*(6); in this case all polygon diagonals are of unit length.

For unambiguity, we will consider all *LSP*(*n*) instances with a fixed position corresponding to appropriate versions of Figure 1 for even values $n \geq 6$. Following the standard postulated assumption, each *n*-polygon considered here is symmetrical with respect to the diameter that connects its "lowest" positioned vertex (that can be placed at the origin) with its "highest" vertex.

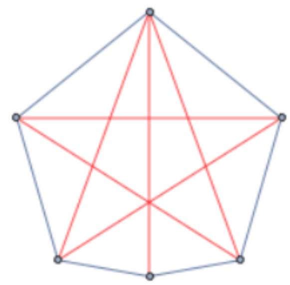

Figure 1. $LSP(6)$

Source: http://mathworld.wolfram.com/GrahamsBiggestLittleHexagon.html

Reinhardt (1922) proved that for all *odd* values $n \geq 3$, $LSP(n)$ is the regular n-polygon. Perhaps surprisingly, this statement is not valid for *even* values of n. For $n = 4$, the square with diameter 1 has maximum area, but an infinite number of other tetragons have the same area. The case $n = 6$ (hexagon) was analyzed and solved by Graham (1975); similarly, the case $n = 8$ (octagon) was solved by Audet *et al.* (2002). Recently, Henrion and Messine (2018) found the largest small polygons for $n = 10$ (decagon) and $n = 12$ (dodecagon), and presented rigorous bounds for the optimum value for $n \leq 16$. We refer to these studies and their topical references for theoretical background and further details regarding the rigorous analysis of the $\{LSP(n)\}$ problem class. We also review the results obtained by using general-purpose local nonlinear optimization software, as reported by Bondarenko *et al.* (1998), Dolan and Moré (2000).

Here we follow a numerical global optimization approach, in order to find $LSP(n)$ configurations and corresponding conjectured $A(n)$ values. Following this introduction, one of the standard optimization model forms is presented in section 2. Earlier alternative solution approaches and best known results are reviewed in section 3. AMPL and the AMPL-LGO solver option are briefly discussed in section 4, followed by AMPL-LGO results (section 5), regression model development and resulting conjectured optimum values $\{A(n)\}$ for $n \geq 6$ (section 6). Conclusions are presented in section 7.

2 Optimization Model

Our objective is to find numerically optimized $LSP(n)$ instances with $n \geq 6$ vertices. The model formulation presented here is cited from Bondarenko *et al.* (1998) who refer to Gay's model (1998), discussed also by Gay (2015). His model code, named pgon.mod, refers to a GAMS model developed by Francisco J. Prieto (a more precise reference to Prieto's work is unknown to this author). The corresponding GAMS (2018) model library item polygon.gms refers back to Gay, to Graham's article, and to the benchmarking study by Dolan and Moré (2000).

Following the model formulations referred to above, we use polar coordinates to describe the n-polygon, assuming that vertex i is positioned at polar radius r_i and at angle θ_i. For unambiguity, we assume that the vertices are arranged (indexed) according to increasing angles. Placing the last vertex position at the origin, we set $r_n = 0$, $\theta_n = \pi$. (We refer again to Figure 1 for the hexagon instance *LSP*(6) that corresponds to this standardized position.)

LSP(*n*) Model Formulation

Maximize total area of the *n*-polygon:

$Max\ A(n) = ½ \sum_{i=1,\ldots,n-1} r_i\, r_{i+1} \sin(\theta_{i+1} - \theta_i)$.

Constraints related to the prescribed pairwise distance of vertices *i* and *j*:

$r_i^2 + r_j^2 - 2\, r_i\, r_j \cos(\theta_i - \theta_j) \leq 1$, for $1 \leq i \leq n - 2$, $i + 1 \leq j \leq n - 1$.

Vertex angle ordering relations:

$\theta_{i+1} - \theta_i \geq 0$, for $1 \leq i \leq n - 2$.

Variable bounds and fixed settings:

$0 \leq \theta_i \leq \pi$ and $0 \leq r_i \leq 1$, for $1 \leq i \leq n - 1$; $r_n = 0$, $\theta_n = \pi$.

Numerical challenges are expected to arise due to the nonconvexity of the objective function and of the nonlinear constraints: the number of nonlinear constraints increases *quadratically* as a function of *n*. For example, the *LSP*(80) model instance has 158 decision variables with bound constraints (given the fixed values for r_n and θ_n); and it has 3241 constraints of which 3161 are nonconvex (adding to the 78 linear constraints the two fixed value constraints).

As noted by other researchers and numerically supported also by the present study, while the standardized *LSP*(*n*) model instances have a unique global solution, the number of local optima increases with *n*. Many of the local optima are close in quality to the (unknown or only approximately known) global optimum: this fact makes the *LSP*(*n*) problem-class numerically challenging – similarly to many other object configuration design problems arising e.g. in computational physics, chemistry and biology.

3 Related Earlier Studies and Results

3.1. Analytical Approaches

Following Graham (1975) – who applies geometric insight and results by Woodall (1971) – finding *LSP*(6) requires the *exact* solution of a 10^{th} order irreducible polynomial equation. More specifically, the area *A*(6) of *LSP*(6) can be found as the second-largest real root *r* of the equation

$$11993 - 78488\, r + 144464\, r^2 + 1232\, r^3 - 221360\, r^4 + 146496\, r^5 + 21056\, r^6 - 30848\, r^7 - 3008\, r^8 + 8192\, r^9 + 4096\, r^{10} = 0$$

It can be expected that this and similar calculations become hard and numerically intensive – even in the most sophisticated computational environments of today. The reader can verify this remark by evaluating the *Mathematica* notebook called GrahamsLargestSmallHexagon, downloadable from Weisstein's Biggest Little Hexagon webpage.

Audet *et al.* (2002), Henrion and Messine (2018) follow a different approach: in their studies finding *LSP*(*n*) requires the *exact* solution of a corresponding nonconvex quadratic programming problem with quadratic constraints, combined again with geometric analysis. This approach (based on a different model from the one cited in section 2) also brings the LSP problem-class into the realm of global optimization.

Henrion and Messine (2018) conjecture that *LSP*(*n*) for all even values $n \geq 4$ has a symmetry axis, as indicated by Figure 1 for *LSP*(6). This conjecture was proved by Reinhardt (1922) for $n = 4$, and by Yuan (2004) for $n = 6$. As noted by Henrion and Messine, Graham used this conjecture to find *LSP*(6); the *LSP*(8) configuration found by Audet *et al.* (2002) also supports the conjecture. The software packages SeDuMi (Storm, 1999), VSDP (Jansson, 2006), and GloptiPoly (Henrion *et al.*, 2009) are used by Henrion and Messine to solve *LSP*(*n*) instances for $6 \leq n \leq 16$. Henrion and Messine (2018) discuss the current computational limitations of this approach, as runtimes rapidly grow from seconds to tens of minutes in their numerical tests.

Table 1 summarizes *all* currently known validated numerical results, including also the validated bounds found by Henrion and Messine (2018).

Table 1. Numerical results based on analytical approaches

n	*LSP*(*n*) area *A*(*n*)	References
4	0.5	Reinhardt (1922)
6	0.674981..	Graham (1975)*
8	0.726867..	Audet *et al.* (2002)*
8	$0.72686845 \leq A(8) \leq 0.72686849$	Henrion and Messine (2018)**
10	$0.74913721 \leq A(10) \leq 0.74913736$	
12	$0.76072986 \leq A(12) \leq 0.76072997$	
14	$0.76753100 \leq A(14) \leq 0.76893595$	
16	$0.77185969 \leq A(16) \leq 0.77279135$	

* The numerical results given with 6-digit precision are cited from Graham (1975), Audet *et al.* (2002), and from Weisstein's webpage titled Biggest Little Polygon: consult the related references.
** Notice the slight numerical discrepancy between the results of Audet *et al.* (2002) and Henrion and Messine (2018) for the case $n = 8$.

3.2 Numerical Solution Approaches

The COPS technical report by Bondarenko *et al.* (1998) presents comparative numerical results for several *LSP*(*n*) instances as shown in Table 2. These results were obtained by using the *local* nonlinear optimization software packages DONLP2, LANCELOT, LOQO, MINOS, and SNOPT (September 1998 versions) linked to the AMPL modeling environment. Table 2 summarizes the *best* numerical solution – obtained by at least one of the listed solvers – cited from the COPS report. The term *best* refers to the solution which has the highest objective function value, while meeting all model constraints with at least 10^{-8} precision. For completeness, we also added results for n = 25, 50, 75, 100 from the subsequent benchmarking study of Dolan and Moré (2000). They tested LANCELOT, LOQO, MINOS, and SNOPT: again, we cite only the overall *best* results.

Table 2. Numerical results obtained by general-purpose *local* nonlinear optimization software Bondarenko *et al.* (1998), Dolan and Moré (2000)

Best results found by at least one of DONLP2, LANCELOT, LOQO, MINOS, SNOPT

n	$LSP(n)$ area $A(n)$
6	0.6749814429
10	0.7491373458
20	0.7768587560
25	0.779740..
50	0.7840161480
75	0.784769..
100	0.7850565708

* The best numerical results given with 6-digit precision are cited from Dolan and Moré (2000).

To the author's knowledge, the results presented in Tables 1 and 2 encompass *all* publicly available results for the entire $\{LSP(n)\}$ problem-class.

3.3 Additional Remarks

Following the model cited in section 2, the numerical solution of $LSP(n)$ instances requires the handling of a nonlinear programming problem with $O(n^2)$ nonconvex constraints, and a nonconvex objective function. In spite of this perceived difficulty, the $LSP(n)$ problems are thought not to become "dramatically" more difficult to handle as n increases. This opinion is based on the conjectured structural similarity and symmetry of $\{LSP(n)\}$ configurations. As noted also in the COPS report, it is expected that the optimal $LSP(n)$ configurations approach the circle of unit diameter as $n \to \infty$. Therefore we can conjecture the limit relation $A(\infty) = \pi/4 \sim 0.7853981634$. However, the *numerical* challenge – due to the presence of an increasing number of local optima with values close to the global optimum – still remains.

Based on the insight mentioned above, the COPS report uses as initial solution guess "a polygon with almost equal sides" (without further details presented in the report, but demonstrated by the AMPL and GAMS code implementations referred to earlier). This remark illustrates the point that, in optimization models which have a "plausible" structure, good insight can be an essential step towards finding credible solutions efficiently – however, without provable global optimality in many cases. In spite of a "hand-crafted" initial solution, the high-quality local solvers listed above sometimes failed to find solutions, and the solutions returned were typically somewhat different and often evidently suboptimal. For further details, consult Bondarenko *et al.* (1998), Dolan and Moré (2000), and the results presented later on in Tables 3 and 4.

Here we follow a numerical optimization approach, keeping in mind also the above cautionary notes. Specifically, using the LGO global-local optimization solver engine linked to the AMPL modeling environment, we present *conjectured* numerical results for all even values $6 \le n \le 80$.

Our results are in close agreement with the best results reported in Tables 1 and 2. For comparison, we also report results attained by the solvers MINOS, SNOPT (until reaching the set demo license size limitations), and IPOPT (a freely available solver linked also to AMPL). In this study, we refer only to the IPOPT (2018), MINOS (1998) and SNOPT (2006) documentation, available with the current AMPL implementation. References to all other solver engines mentioned, but not used here are available from their respective documentation.

4 AMPL and AMPL-LGO

4.1 AMPL

AMPL is a powerful modeling language that facilitates the formulation of optimization models and the generation of the requisite computational data structures. AMPL enables model development in a natural, concise, and scalable fashion; it also supports the seamless invocation of suitable solver engines to handle the optimization models created. AMPL has been extensively documented elsewhere: therefore here we refer only to the AMPL book by Fourer *et al.* (2003), and to the extensive resources available at the AMPL website www.ampl.com.

4.2. LGO

Nonlinear optimization models frequently have multiple – local and global – optima: the objective of global optimization is to find the best possible solution under such circumstances. LGO is an integrated global-local solver suite for constrained nonlinear optimization. The model-class addressed by LGO is concisely defined by the vector of decision variables $x \in R^n$; the explicit, finite n-vector variable bounds l and u; the continuous objective function $f(x)$; and the (possibly absent) m-vector of continuous constraint functions $g(x)$. Applying these notations, LGO is aimed at solving models of the form

$$(1) \qquad \min f(x) \quad \text{subject to} \quad x \in D := \{x: l \leq x \leq u, g(x) \leq 0\}.$$

Obviously, in (1) all vector inequalities are interpreted component-wise (l, x, u, are n-component vectors and 0 denotes the m-component zero vector). Formally more general optimization models that include also = and ≥ constraint relations and/or explicit lower and upper bounds on the constraint function values can be simply reduced to the model form (1). If D is non-empty, then the stated basic analytical assumptions guarantee that the optimal solution set X^* of the model is non-empty, while finding X^* could still be a formidable analytical and/or numerical challenge. Clearly, the $\{LSP(n)\}$ problem-class is encompassed by the generic optimization model (1).

Without going into further details here, let us mention that the theoretical foundations of the LGO software development project are discussed by Pintér (1996); computational implementation aspects are discussed e.g., by Pintér (2002, 2009). Here we utilize the LGO solver option available for use with AMPL (Pintér, 2015); the current stand-alone LGO implementation is documented in Pintér (2017). In addition to these references, e.g., the recent studies by Pintér (2018) and Pintér *et al.* (2018) present numerical results using LGO to solve a range of nonlinear optimization problems, from relatively simple standard test problems to well-known challenges.

5 Numerical Results and Comparisons

5.1 AMPL-LGO Results

In our numerical tests, the AMPL code implementation pgon.mod was used. All test runs were conducted on a (several years old) laptop PC with Intel Core i5-3337-U CPU @ 1.80 GHz (x-64 processor), 16 Gb RAM, running under the Windows 10 (64-bit) operating system.

The results of a *single, completely reproducible* run-sequence are summarized in Table 3 for all even values $6 \le n \le 80$, with a single setting of LGO solver options. In several – seemingly more difficult – cases, we received somewhat better numerical results in additional tests, at the expense of longer runtimes: for consistency we did not include those results here. All $A(n)$ values are directly cited from the AMPL-LGO solver output; the corresponding $LSP(n)$ configurations are automatically written to a result text file. To avoid reporting excessive details, the configurations found are not presented: all can be reproduced (as needed) in a matter of seconds, for each case $6 \le n \le 80$ considered here.

Table 3. AMPL-LGO numerical results

n	*LSP*(*n*) area *A*(*n*)	Runtime (seconds)	Maximum constraint violation
6	0.6749814433	0.55	2.21e-09 (i.e., $2.21 \cdot 10^{-9}$, etc.)
8	0.7268684830	0.70	6.47e-09
10	0.7491373457	0.95	2.96e-10
12	0.7607298709	1.30	8.9e-10
14	0.7675310106	1.69	3.91e-09
16	0.7718613224	2.55	4.09e-09
18	0.7747881650	2.63	9.78e-09
20	0.7768587506	3.02	2.23e-09
22	0.7783773308	3.95	9.08e-09
24	0.7795240461	5.22	7.73e-09
26	0.7804111201	5.34	6.34e-09
28	0.7811114192	6.05	9.83e-09
30	0.7816739255	6.98	3.67e-09
32	0.7818946320	5.72	6.29e-10
34	0.7823103007	7.61	9.03e-09
36	0.7826513767	9.50	9.75e-09
38	0.7829526627	9.34	5.08e-09
40	0.7832011589	9.55	8.47e-11
42	0.7834135187	12.06	4.62e-09
44	0.7835966860	13.22	1.42e-09
46	0.7837554636	16.88	3.43e-09
48	0.7838942710	17.95	8.31e-09
50	0.7840161496	16.53	9.99e-09
52	0.7841233641	20.61	8.78e-09

54	0.7842192995	21.38	9.18e-09
56	0.7843044654	23.91	3.87e-09
58	0.7843807534	22.95	8.43e-09
60	0.7844492943	27.97	9.79e-09
62	0.7845111362	21.22	8.93e-09
64	0.7834620877	30.48	9.82e-09
66	0.7845910589	34.17	1.19e-09
68	0.7846139029	35.84	9.00e-09
70	0.7846403575	22.33	6.45e-09
72	0.7847454020	42.72	7.34e-09
74	0.7845564840	26.25	3.54e-09
76	0.7847585719	49.19	8.95e-09
78	0.7845160579	49.47	9.64e-09
80	0.7848252941	51.45	7.25e-09

Our numerical results are in fairly close agreement with all (best known at the time of publication) results displayed in Tables 1 and 2, with small discrepancies. In several cases, we have found somewhat better *conjectured* optimum estimates compared to the earlier results (which are also numerical estimates). The runtimes appear to scale well for $6 \leq n \leq 80$, mostly (but not always) increasing with *n*. The entire sequence of the 38 optimization runs reported here took a little over 10 minutes.

Although AMPL-LGO seems to perform fairly well in comparison to the other solvers tested by us or by others, its numerical limitations start to show at (or around) $n = 64$. The results presented in Table 3 for $n = 64$, 74, and 78 are clearly at least somewhat suboptimal, while all other $A(n)$ values are monotonically increasing with *n*, as expected. Instead of "tweaking" the LGO option parameters – e.g., by increasing the global search effort limit (which was actually reached in several cases reported above, for some of the larger *n* values), or increasing the time limit (set to 5 minutes for each run, and never reached) – here we simply apply linear interpolation to "rectify" the suboptimal results based on the "bracketing" values in Table 3. For example, $A(64)$ is estimated on the basis of the results obtained for $A(62)$ and $A(66)$. This interpolation leads to the following estimated values:

$A(64) \sim 0.7845510976$, $A(74) \sim 0.7847519869$, $A(78) \sim 0.7847919330$.

These estimated values will be used in developing a regression model in section 6.

5.2. An Illustrative Comparison with Other AMPL Solver Results

For a somewhat more comprehensive picture, we also generated a set of comparative results using several currently shipped AMPL solvers, see Table 4. (The LGO results are cited from Table 3.)

Table 4. Comparative numerical results obtained by several current AMPL solvers

n	*LSP*(*n*) area A(*n*)			
	MINOS	SNOPT	IPOPT	LGO

6	0.6749814429	0.6749814429	0.6749814308	0.6749814433
8	0.7268684828	0.7268684827	0.7268684678	0.7268684830
10	0.7491373459	0.7491373459	0.7371215901	0.7491373457
12	0.7607298734	0.7607298734	0.7542668597	0.7607298709
14	0.7521931121	0.7675310112	0.7675309793	0.7675310106
16	0.7625954979	0.7718613220	0.7696844715	0.7718613224
18	0.7554106917	0.7747881651	0.7491373424	0.7747881650
20	0.7649920891	0.7768587560	0.7732071277	0.7768587506
22	0.7640946468	0.7783773301	0.7607298336	0.7783773308
24	0.7640946468	0.7795240452	0.7548403603	0.7795240461
26	* demo	* demo	0.7523851367	0.7804111201
28			0.7523851373	0.7811114192
30			0.7491373081	0.7816739255
40			0.7268684622	0.7832011589
50			0.7197409051	0.7840161496
60			0.6749814462	0.7844492943
70			0.7268685003	0.7846403575
80			0.7197409068	0.7848252941

* demo license (solvable model size) limit attained

As we can see, the other solvers return close, but somewhat different results for small values of *n*. The numerical limitations start to become more apparent as *n* increases.

6 Regression Model and Conjectured Optima

Based on the numerical results obtained by AMPL-LGO, next we present a nonlinear regression model that enables the estimation of the optimal area A(*n*), for arbitrary (even) values of *n*. This regression model has been determined using the NonlinearModelFit function of *Mathematica* (Wolfram Research, 2018).

Given that $A(n)$ is a monotonically increasing function of *n*, and $A(\infty) = \pi/4$, the following model form is proposed:

(2) $$A(n) = \pi/4 - c_1/n - c_2/n^2 - c_3/n^3.$$

Using the estimated A(*n*) values found by AMPL-LGO, except substituting the three suboptimal A(*n*) values by their interpolated approximation, the following regression model is found (rounding the coefficients found to five digits after the decimal point):

(3) $$A(n) \sim \pi/4 - 0.01098/n - 2.91512/n^2 - 5.96369/n^3.$$

Applying this model, we obtain the illustrative $A(n)$ estimates shown in Table 5; all rounded to 6-digit precision after the decimal point, in line with the least precise results cited earlier.

Table 5 Estimated $A(n)$ values based on the regression model (3)

n	6	8	10	20	30	40
A(*n*)	0.674983	0.726829	0.749185	0.776816	0.781572	0.783209
n	50	60	70	80	90	100
A(*n*)	0.783965	0.784378	0.784629	0.784794	0.784908	0.784991
n	200	300	400	500	1000	2000
A(*n*)	0.785270	0.785329	0.785352	0.785364	0.785384	0.785392

All values are in reasonably close agreement with the best numerical results presented earlier in Tables 1 to 4 (when such values are available, up to $n = 100$); and provide additional optimum estimates for an illustrative selection of values $n > 100$. Since the data used to develop the regression model (2) are likely to be at least slightly suboptimal, one can expect that the estimated $A(n)$ values will become also (slightly) suboptimal.

For a more complete analysis and for comparative purposes, first and second order regression models (with $1/n$ as their input argument) were also calculated, but the third order model (2) with values in (3) clearly resulted in a superior fit to the entire data set used. Obviously, within reason higher order models (or perhaps other model types) could give even more precise fit to the data, but – considering also the inherent data inaccuracies – the third order model already gives a fairly good fit. Figure 2 shows the model function curve defined by (3) together with the revised data set (dots) that includes the interpolated data.

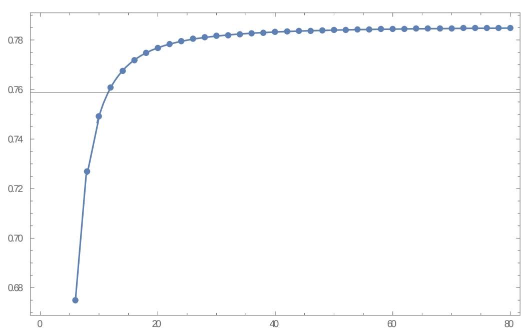

Figure 2.

Nonlinear regression model (3) *vs.* the revised data set for $6 \leq n \leq 80$.

Figure 3 displays the regression model residuals. With a few exceptions, the residual errors are less than $1 \cdot 10^{-4}$; the absolute value of the largest estimated error is ~ $8 \cdot 10^{-4}$. These estimated error values are reasonably small, compared to the range of the observed data ~ [0.674981, 0.784825].

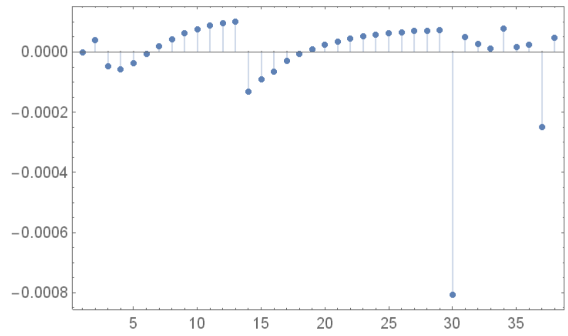

Figure 3.

Residuals in the regression model of $A(n)$, $6 \leq n \leq 80$.

Notice that most of the residuals seem to follow an interesting cyclical pattern that – in the author's opinion – seems more due to the inherent structure of the $LSP(n)$ problem-class than to numerical fluctuations and other "noise" induced by the computational environment used.

7 Conclusions

In this article we study the problem of finding the sequence of largest small n-polygons $LSP(n)$ with unit diameter and maximal area $A(n)$. Finding $LSP(n)$ and $A(n)$ for even values of $n \geq 6$ is a long-standing challenge, leading to an interesting class of nonlinear optimization problems.

The structural properties of this problem, and of similar optimization challenges – notably: atomic structure models, potential energy models, regular object packings, and many other (identical object configuration based) optimization problems – often support the finding of "credible" initial solutions and solution guesses. However, finding the true global solution remains difficult, as our numerical study also illustrates.

Using the AMPL modeling environment with the LGO solver option, we present global search based numerical solutions for all even values $6 \leq n \leq 80$. Our results are comparable to (in a number of cases are better than) the results obtained by other authors and other solver software. Based on the results obtained, we propose a regression model that enables the estimation of the optimal area sequence $\{A(n)\}$, for arbitrary integer values of n.

The motto of the COPS benchmarking studies cited here (Bondarenko *et al.*, 1998; Dolan and Moré, 2000) is arguably somewhat funny, but the message is worth quoting: "COPS: Keeping optimization software honest." (See https://www.mcs.anl.gov/~more/cops/.) In accordance with this message, we wish to add some honest as well as pragmatic advice, not driven by unconditional "software developer's pride". Facing the vast universe of nonlinear optimization problems, it is wise to refrain from overly confident statements regarding the superiority of any particular solver software over the others. Instead, it is salient practice to use a repertoire of appropriate solver options whenever possible, since it may not be obvious *a priori* which solver option will work best for a specific – new, or unusually hard – model.

To end on an admittedly cheerful-but-serious note, two or more sensible heads are frequently better then just one. This fact has been observed not only by honest researchers, but also by many others with an artistic or philosophical leaning. For sophisticated melody and a related seriously funny message, consult – and enjoy – Gray and Ross (1949).

Acknowledgements

The author thanks David M. Gay (AMPL Optimization, Inc.) for his contributions to the AMPL-LGO solver link development, and for useful discussions. Thanks are due also to Wolfram Research for their long-standing support of the author's research.

References

AMPL Optimization (2018) AMPL. www.ampl.com.

Audet, C., Hansen, P., Messine, F., and Xiong, J. (2002) The Largest Small Octagon. *Journal of Combinatorial Theory*, Series A, 98: 46-59.

Bondarenko, A.S., Bortz, D.M., Moré, J.J. (1998) COPS: Large-Scale Nonlinearly Constrained Optimization Problems. Technical Report ANL/MCS-TM-237, Argonne National Laboratory, Argonne, IL.

Dolan, E.D. and Moré, J.J. (2000) Benchmarking Optimization Software with COPS. Technical Report ANL/MCS-TM-246, Argonne National Laboratory, Argonne, IL.

Fourer, R. Gay, D.M. and Kernighan, B.W. (2003) *AMPL: A Modeling Language for Mathematical Programming* (2nd Edition). Brooks/Cole-Thomson Learning, Pacific Grove, CA. Available also at www.ampl.com/resources/the-ampl-book/.

GAMS (2018) Model libraries; polygon.gms : Largest small polygon COPS 2.0 #1. https://www.gams.com/latest/gamslib_ml/libhtml/gamslib_polygon.html.

Gay, D.M., (1998) http://netlib.org/ampl/models/pgon.mod.

Gay, D.M. (2015) The AMPL Modeling Language: An Aid to Formulating and Solving Optimization Problems. In: Al-Baali, M., Grandinetti, L., and Purnama, A., Eds., *Numerical Analysis and Optimization: NAO-III (Muscat, Oman, January 2014)*, pp. 95-116. Springer International Publishing, Switzerland.

Gill, P.E., Murray, W., and Saunders, M.A. (2006) User's Guide for SNOPT Version 7: Software for Large-Scale Nonlinear Programming. Systems Optimization Laboratory, Department of Management Science and Engineering, Stanford University, Stanford, CA.

Graham, R.L. (1975) The Largest Small Hexagon. *Journal of Combinatorial Theory*, Series A, 18: 165–170.

Gray, W. and Ross, A. (1949) Twisted. Composed by Wardell Gray, lyrics by Annie Ross. (Performed by numerous artists, including Joni Mitchell on the album *Court and Spark*, 1974.)

Henrion, D. Lasserre, J. B. (2003) GloptiPoly: Global Optimization over Polynomials with Matlab and SeDuMi. *ACM Transactions on Mathematical Software*, 29 (2), 165–194.

Henrion, D., Lasserre, J.B., and Loefberg, J. (2009) GloptiPoly 3: Moments, Optimization and Semidefinite Programming. *Optimization Methods and Software* 24 (4-5), 761-779.

Henrion, D. and Messine, F. (2018) Finding Largest Small Polygons with GloptiPoly. Available at https://arxiv.org/pdf/1103.4456.pdf (Current version: June 17, 2018).

IPOPT (2018) Ipopt, a Library for Large-scale Nonlinear Optimization. http://projects.coin-or.org/Ipopt. The authors of Ipopt are listed at https://projects.coin-or.org/Ipopt/browser/trunk/Ipopt/AUTHORS.

Jansson, C. (2006) VSDP: A Matlab Software Package for Verified Semidefinite Programming. *Nonlinear Theory and its Applications*, 327–330.

Murtagh, B.A. and Saunders, M.A. (1983…1998) MINOS 5.5 User's Guide. Technical Report SOL 83-20R. Systems Optimization Laboratory, Department of Management Science and Engineering, Stanford University, Stanford, CA. (Revised version: July 1998)

Pintér, J.D. (1996) *Global Optimization in Action*. Kluwer Academic Publishers, Dordrecht.

Pintér J.D. (2002) Global Optimization: Software, Test Problems, and Applications. In: Pardalos, P.M. and Romeijn, H.E., Eds. *Handbook of Global Optimization, Volume 2*, pp. 515-569. Kluwer Academic Publishers, Dordrecht.

Pintér, J.D. (2009) Software Development for Global Optimization. In: Pardalos, P.M. and T. F. Coleman, Eds. *Global Optimization: Methods and Applications*, pp. 183-204. Fields Institute Communications Volume 55. Published by the American Mathematical Society, Providence, RI.

Pintér, J.D. (2015) *AMPL-LGO User's Guide*. Available from AMPL Optimization, Inc. and from the author.

Pintér, J.D. (2017) *LGO Solver Suite for Global and Local Nonlinear Optimization. User's Guide.* Available from the author. PCS Inc., www.pinterconsulting.com.

Pintér, J.D. (2018) How Difficult is Nonlinear Optimization? A Practical Solver Tuning Approach, with Illustrative Results. *Annals of Operations Research* 265, 119–141.

Pintér, J.D., Kampas, F.J., Castillo, I. (2018) Globally Optimized Packings of Non-uniform Size Spheres in R^d: A Computational Study. *Optimization Letters* 12 (3), 585–613.

Reinhardt, K. (1922) Extremale Polygone gegebenen Durchmessers, *Jahresbericht der Deutschen Mathematiker-Vereinigung* 31, 251–270.

Sturm, J. F. (1999) Using SeDuMi 1.02, A Matlab Toolbox for Optimization over Symmetric Cones. *Optimization Methods and Software* 11 (1-4), 625–653.

Weisstein, E.W., "Biggest Little Polygon." From MathWorld – A Wolfram Web Resource. http://mathworld.wolfram.com/BiggestLittlePolygon.html.

Weisstein, E.W. "Graham's Biggest Little Hexagon." From MathWorld – A Wolfram Web Resource. http://mathworld.wolfram.com/GrahamsBiggestLittleHexagon.html.

Wolfram Research (2018) *Mathematica* (version 11.0). Wolfram Research, Inc., Champaign, IL.

Woodall, D.R. (1971) Thrackles and Deadlock. In: Welsh, D. J. A., Ed. *Combinatorial Mathematics and Its Applications*, pp. 335-347. Academic Press, New York.

Yuan, B. (2004) The Largest Small Hexagon. M.Sc. thesis, Department of Mathematics, National University of Singapore.