# Visualization of Nonlinear Programming for Robot Motion Planning

David Hägele University of Stuttgart Stuttgart, Germany david.haegele@visus.uni-stuttgart.de Moataz Abdelaal University of Stuttgart Stuttgart, Germany moataz.abdelaal@visus.unistuttgart.de

Ozgur S. Oguz University of Stuttgart Stuttgart, Germany ozgur.oguz@ipvs.uni-stuttgart.de

Marc Toussaint Technische Universität Berlin Berlin, Germany toussaint@tu-berlin.de Daniel Weiskopf University of Stuttgart Stuttgart, Germany daniel.weiskopf@visus.unistuttgart.de



# Figure 1: User interface of our visual analytics system for nonlinear programming processes. Using multiple coordinated views, a user can assess the evolution of a problem's solution, constraint values, and progression speed throughout an optimization.

## ABSTRACT

Nonlinear programming targets nonlinear optimization with constraints, which is a generic yet complex methodology involving

VINCI 2020, December 8-10, 2020, Eindhoven, Netherlands

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-8750-7/20/12...\$15.00 https://doi.org/10.1145/3430036.3430050 humans for problem modeling and algorithms for problem solving. We address the particularly hard challenge of supporting domain experts in handling, understanding, and trouble-shooting high-dimensional optimization with a large number of constraints. Leveraging visual analytics, users are supported in exploring the computation process of nonlinear constraint optimization. Our system was designed for robot motion planning problems and developed in tight collaboration with domain experts in nonlinear programming and robotics. We report on the experiences from this design study, illustrate the usefulness for relevant example cases, and discuss the extension to visual analytics for nonlinear programming in general.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

## CCS CONCEPTS

• Human-centered computing → Visual analytics; • Mathematics of computing → *Mathematical optimization*; • Computer systems organization → Robotics.

## **KEYWORDS**

Visual analytics, nonlinear programming, trajectory visualization, optimization

### **ACM Reference Format:**

David Hägele, Moataz Abdelaal, Ozgur S. Oguz, Marc Toussaint, and Daniel Weiskopf. 2020. Visualization of Nonlinear Programming for Robot Motion Planning. In *The 13th International Symposium on Visual Information Communication and Interaction (VINCI 2020), December 8–10, 2020, Eindhoven, Netherlands.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/ 3430036.3430050

## **1** INTRODUCTION

Nonlinear constraint optimization, also known as nonlinear programming (NLP), deals with finding optima within constrained sets of variables. Only little work has been published concerning visualization in this field, which is surprising since a substantial amount of research has been conducted in the closely related fields of discrete and also unconstrained optimization. NLP comprises two separate stages: the modeling stage and the solving stage. In the modeling stage, the problem needs to be expressed in terms of an objective function and associated equality or inequality constraint functions that may be nonlinear. A classical example from economics is to optimize spending for greatest profit (objective) while staying within budget (constraint). In the solving stage, an iterative algorithm moves through the (likely high-dimensional) space of the modeled problem to find the optimal location with respect to the objective while making sure to satisfy all of the imposed constraints.

Due to high dimensionality, large number of constraints, and nonlinearity, which gives rise to local optima and disconnected feasible regions, it is challenging to grasp such optimization problems. Even more challenging is the comprehension of unexpected or unsatisfactory behavior of a solver when applied to a problem.

To help NLP experts in understanding their problems better as well as the corresponding behavior of the applied solvers, we want to leverage visualization of the internal steps of an optimization process. We propose a visual analytics approach for a post-mortem analysis of optimization runs of robot motion planning problems that we developed in tight collaboration with domain experts. We focus on visualizing the high-dimensional optimization landscape as seen by the optimizer, to be able to reason about its behavior, and on representing the evolution of the solution throughout the optimization in order to be able to interpret high-dimensional loci as intermediate solutions to the motion planning problem. Our contributions are: 1) a visual analytics system for the analysis of constrained optimizations for robot motion planning, 2) a report of our design study process and lessons learned.

## 2 RELATED WORK

This section is divided into two parts. First, we review related work in the field of visualization of optimization that includes NLP, linear programming, constraint programming, multi-objective optimization, as well as unconstrained optimization. Then, we discuss related work concerned with the visualization of temporal highdimensional data.

## 2.1 Visualization of Optimization

Despite its strong potential [17, 31], the area of visualizing NLP remains unexplored to a large extent. Androulakis and Vrahatis [3] proposed OPTAC, a tool for analyzing and visualizing the convergence behavior of unconstrained optimization algorithms. Charalambos and Izquierdo [9] visualize the geometric shapes of the feasible regions of linear programs. The method uses three-dimensional Cartesian coordinates and therefore, is limited to three-dimensional optimization problems. To display high-dimensional planes, Chatterjee et al. [10] use parallel coordinates plots instead. Since the geometry in linear programs is simple and solvers are fundamentally different, these approaches are not applicable for NLP.

The area of constraint programming, in contrast to NLP, received a lot of attention from visualization practitioners. Most of the work done in this area focuses on visualizing the search tree resulting from constraint programs [8, 17, 34, 36, 37]. However, these techniques cannot be applied to NLP, due to the differences in modeling and solvers [19]. While solving constraint programs involves tree traversal and dynamic programming, NLPs consist of differentiable implicit surfaces, and their solvers are based on algebraic methods such as gradient descent.

Instead of focusing on visualizing the search tree resulting from constraint programming, others attempt to visualize the evolution of variables, constraints, and the interaction between them using matrix views [7, 14, 15]. While matrices provide good representations to explore the relationship between constraints and variables, they have scalability issues, making them only suitable for exploring optimization problems with a small number of variables and constraints. Despite the differences, we find our work shares the same goal, i.e., exploring the evolution of variables and constraints.

The visualization of multi-objective optimization is yet another neighboring field to NLP. Most of the work in this field is concerned with visualizing the solution set in the objective space. Therefore, different visualization methods for high-dimensional data are used. Other methods apply dimensionality reduction to map the highdimensional objective space into a two-dimensional space. Tušar and Filipič [42] provide a comprehensive review of these methods. Although multi-objective optimization and NLP address two different problems, they both share the notion of high dimensionality.

The visualization of high-dimensional unconstrained problem optimization such as in neural networks is discussed by Goodfellow et al. [16], who use a straight line from initialization to found optimum to sample and analyze the loss function. This gave rise to the loss landscape visualization technique by Li et al. [29]. They use a 2D plane to sample the loss function and obtain a contour plot to analyze a subspace into which the optimizer's trajectory can be projected. We extend this technique for NLP to show constraints within the landscape of the objective function. Visualization of Nonlinear Programming for Robot Motion Planning

## 2.2 Visualization of Temporal High-Dimensional Data

There are numerous methods for visualizing high-dimensional data [30], scatterplot matrices [2], glyphs [12], parallel coordinates plots [20], and star coordinates [22]. These techniques, however, do not scale with a large number of dimensions. Therefore, other approaches project high-dimensional data into low-dimensional space using various dimensionality reduction techniques [33]. Introducing the time dimension poses a visualization challenge, as the visual representation needs to convey not only the relation between the different dimensions but also their temporal context and evolution. Aigner et al. [1] provide a comprehensive overview of such time-oriented visualization.

A rather straightforward approach is to include time as an additional dimension in the traditional high-dimensional data representations. For example, Wong and Bergeron add an axis for the time dimension in parallel coordinates [46]. Similarly, TimeWheel [38] arranges the other axes in a circular layout around the time axis. These however, provide a poor representation of the temporal context information and are prone to become cluttered. Other approaches use a depth cue to encode the temporal information, resulting in 3D parallel coordinates [18, 44], 3D star coordinates [32], or space-time cubes [4, 5]. These methods, however, have scalability concerns regarding the size of the data. Additionally, the use of interaction and animation is a necessity to avoid occlusion problems.

More recent approaches use dimensionality reduction to show temporal information. Jäckle et al. [21] proposed temporal multidimensional scaling (TMDS) for visualizing multivariate time series data. Bach et al. [6] and van den Elzen et al. [43] showed the temporal progression of datasets by projecting the individual snapshots of datasets as points in 2D space. The points of subsequent snapshots are then connected by lines, while color is used to encode time.

In our work, we think of the optimization process as a sequence of intermediate solutions. Each solution is characterized by a set of high-dimensional decision variables. Thus, we can adopt the same techniques [6, 43] to visualize the evolution of the optimization process. In our work, in contrast, we deal with two different notions of time simultaneously: the time of a robot's motion and the time of the optimization process.

Torsney-Weir et al. [39] studied multi-dimensional shapes using their hypersliceplorer algorithm. Using separate views for pairs of dimensions does not scale well with our problems, however.

## 3 BACKGROUND

In this section, we summarize the domain background of robot motion planning and nonlinear programming. We also present the requirements for a visualization system that we elicited from the domain experts. From here on we will use *NLP* as an abbreviation for both, nonlinear program and nonlinear programming.

## 3.1 Robot Motion Planning

Motion planning is the problem of finding a collision-free path to move an object from an initial state to a desired goal state [27]. It is a crucial problem in many fields, including robotics, computational biology (drug design, protein folding), virtual prototyping, manufacturing, and computer graphics. In this paper, we focus on robot motion planning as our application domain. However, we expect that our visual analytics approach will carry over to other applications. An example of such motion planning problems is illustrated in Fig. 2.



Figure 2: Animation frames from an optimized robot motion sequence for a box-pushing problem. The robot arm picks up a stick and then moves a box to a target location. The animation was created by our collaborators' motion planning framework.

There are different approaches for finding a valid path connecting the source and target configurations. LaValle [28] provides an overview of the most important ones. Grid- and sampling-based approaches [11, 23–25] discretize the configuration space and use graph search algorithms to find a valid path in the resulting topological space.

Another approach for finding a valid path is based on nonlinear optimization [41]. This involves the formulation and solving of an NLP that describes the desired goal mathematically and imposes constraints to ensure a collision-free path and feasible motion. In contrast to the aforementioned methods, nonlinear optimization does not require any discretization of the configuration space. Therefore, it is possible to find an optimal path if the problem is well-defined.

However, nonlinear optimization requires an initial guess of the solution and the method is prone to get stuck in local minima [13]. To avoid that, it is common to use path finding algorithms as a first step to provide a reasonable initial guess. While this might be enough to find a valid solution, it does not provide insights into the internal mechanism of the optimization algorithms: How fast does the optimizer progress? Or which constraints hinder the algorithm the most? Such questions cannot be answered without a proper investigation of the inner workings of the optimization.

To this end, there is a critical need for visualization tools to debug and troubleshoot the optimization algorithms [17, 31]. Having such tools, can help domain experts formulate hypotheses about the behavior of the optimizer, that could eventually lead to a better reformulation of the cost function and/or the constraints. In this context, we were approached by robotics researchers. During the course of this design study, we have been collaborating with them for more than 9 months, integrating their expertise and knowledge about their domain-specific problems. To better understand the domain problem we tried to characterize it by eliciting requirements for a visual analytics system and tasks to be achievable from the experts. We were able to identify two high-level information needs:

- Q1: When an optimization requires a lot of time to converge, what is taking the optimizer so long?
- Q2: When an optimization produces an infeasible solution, what prevents the optimizer from getting to a feasible location?

To be able to analyze an optimization with regard to these questions, we formulated the following analysis tasks that the experts want to perform on optimizations of their robot motion planning framework:

- T1: Identify and characterize different phases of the optimization process.
- T2: Identify when the optimizer converges faster or more slowly to a local minimum during the optimization process.
- T3: Characterize the evolution of the constraint function values during the optimization process.
- T4: Identify the shape and boundaries of a constraint.
- T5: Identify the shape and boundaries of the feasible regions.

Throughout the project we followed a design study process [35] regularly meeting with our collaboration partners to continuously refine our design and to update requirements.

## 3.2 Nonlinear Programming

Nonlinear programs consist of three parts that allow expressing an optimization problem with auxiliary conditions on the solution, known as the constraints. It typically reads like this:

minimize 
$$f(x)$$
 subject to  $h_i(x) = 0$ ,  $g_j(x) \le 0$  (1)

In this, x is the set of decision variables that can be varied. It makes sense, in our case, to think about x as a large vector describing a robot's motion. The objective function f is to be minimized while all of the inequality constraints  $g_j$  and equality constraints  $h_i$  have to be satisfied at the solution  $x^*$ . In an NLP, at least one of these functions is nonlinear. In high-dimensional space, the equalities define hypersurfaces, and inequalities define hypersurface-confined sets in which the optimal  $x^*$  has to be located.

Including time into an optimization problem is often done by discretization into several time steps. This also applies to our case of robot motion, where the motion path consists of joint configurations representing the robot's poses over time.

The objective and constraints are used to express the robot's task mathematically and to guarantee that the resulting motion is physically viable. In such time-dependent optimization problems, some constraints will typically apply to individual time steps only to enforce valid states, and others will take several times steps into account to enforce valid state transitions or global properties.

This kind of modeling leads to high-dimensional problem spaces and large numbers of constraints. We are dealing with  $\approx 100$  time discrete joint configurations that range from 10 to 30 dimensions, amounting in an optimization space of 1000- or more dimensions, and more than 100 constraints.

*Technical Background:* We now discuss a class of iterative algorithms to solve NLPs. We will not go into details but sketch the general mechanic that is common in the log-barrier, squared penalty, and also the augmented Lagrangian methods. All of these want to find a solution that satisfies the Karush-Kuhn-Tucker conditions [26] (a set of conditions that hold at a feasible minimum of an NLP), and thus determine the dual variables. This is done by constructing an unconstrained problem  $L_{\lambda,\kappa}(x)$  from the NLP that can be minimized using standard techniques like gradient descent or Newtons's method. For example in the augmented Lagrangian method  $L_{\lambda,\kappa}(x) = f(x) + \kappa^{\top}h(x) + \lambda^{\top}g(x) + ||h(x)||_2^2 + \sum_j [g_j(x) > 0]g_j(x)^2$ . The unconstrained problem is then minimized repeatedly while updating the dual parameters  $\lambda, \kappa$ , and thus changing the impact of constraints, in between.

A pseudo code implementation is shown in Listing 1. A line search mechanism (most inner loop decreasing step size) is leveraged to ensure sufficient decrease in each step to prevent overshooting valleys and enforces satisfaction of the Wolfe condition [45].

#### Listing 1: Pseudo Code for NLP Solver

```
FUNCTION solveNLP
 1
 2
    INPUT: objective f, equalities h, and inequalities q
    \ensuremath{\textbf{OUTPUT}} : optimal location x
 3
 4
     BEGIN
      define unconstrained problem L_{\lambda,\kappa} from f , h , and g
 5
 6
      initialize x randomly (or informed)
 7
      DO
       x = argmin of L_{\lambda,\kappa} with initialization x
 8
 9
       update dual parameters \lambda,\kappa
10
      UNTIL convergence
11
     FND
12
13
     FUNCTION argmin
     INPUT: loss function L, initialization x
14
15
     OUTPUT: optimal location x
16
     BEGIN
17
      initialize stepsize a = 1 (or previous value)
18
      WHILE a reasonably large
19
       set descent direction \delta = Newton step for L at x
       WHILE L(x + a\delta) NOT sufficiently smaller than L(x)
20
21
        decrease stepsize a
22
       END WHILE
23
       x = x + a\delta
24
       increase stepsize a
25
      END WHILE
26
    END
```

Besides the problem formulation, i.e., functions f, g, and h, this algorithm is the source of our data for visualization. We define the optimization trajectory as the sequence S of arguments  $x' = x + a\delta$  that are tested during line search in the argmin procedure of Listing 1 (line 20):

$$S = \{x_{\text{init}}, x_1, x_2, .., x^*\}$$
(2)

Furthermore, we obtain a detailed log of the sequence in which different stages in the algorithm were executed: updates to the dual variables (line 9), evaluations of the loss function and corresponding function values of objective and constraints (line 20), step size decrease during line search (line 21), and updates to x (line 23).

## 4 METHOD

In this section, we will introduce our visual analytics system and discuss the special techniques used to visualize the optimization process, i.e., the loss landscape visualization and the robot path evolution visualization.

### 4.1 Visual Analytics System

Our visual analytics system combines and links different views to facilitate the exploration of an optimization process. Due to the iterative nature of the algorithm, the optimization process describes Visualization of Nonlinear Programming for Robot Motion Planning



(b) per configuration trajectories

Figure 3: Figures show the motion path evolution over the optimization process using PCA for projection. (a) It can be observed how the path evolves from a chaotic random shape (top) to a smooth curve (bottom). (b) The trajectory from initialization to final solution for each configuration  $c_t$  is shown and color-coded by time t. The underlying optimization corresponds to an NLP for pushing a box in a circular motion similar to the illustration in Fig. 2.

the evolution of essentially two things: The evolution of objective and constraint function values, and the evolution of the solution described by the sequence of intermediate solutions  $s_i \in S$  on the optimization trajectory.

To show the evolution of function values, which allows for task T3 to be accomplished, we provide line charts for equality and inequality constraints, plotting  $h(s_i)$  and  $g(s_i)$  with optimization step *i* on the *x*-axis (views C and D in Fig. 1). We group constraints by names due to their large number, which allows us to prevent visual clutter. For each group  $H_k$ ,  $G_k$ , we plot the maximum value of all its constraints as aggregation.

$$\bar{h}_k(s_i) = \max_{h_j \in H_k} (h_j(s_i)) \; ; \; \bar{g}_k(s_i) = \max_{g_j \in G_k} (|g_j(s_i)|)$$

Groups are listed in a separate view (F) and can be expanded if desired, to reveal the individual unaggregated constraints in the plot.

To show the evolution of the problem's solution we employ a time curves [6] approach in which we reduce dimensionality of the joint configurations of  $s = \{c_1 \dots c_T\}$ , that is the motion path (view A in Fig. 1). In this 2D representation of the solution, we can show its evolution. Connecting subsequent joint configurations of s through line segments, yields a time curve describing the robot motion. Connecting the same joint configuration of subsequent intermediate solutions  $s_i$ , results in a time curve describing the evolution of that particular configuration throughout the optimization. Figure 3 illustrates this technique that is integrated in view A in Fig. 1.

To combine constraints and solution evolution we adapt the loss landscape technique [29] to NLP, which also enables us to observe the behavior of the optimizer as it travels through high-dimensional space (view B in Fig. 1). Using isobands, we can show contours of the objective function and constraints on a 2D plane slice. The slice plane is spanned by the vector connecting the point of the currently selected optimization step with the minimum  $x^*$ , and a perpendicular vector that is a linear combination of the first two

principal vectors of the optimization trajectory S. Selecting three optimization steps aligns the slice to coincide with the corresponding points. We draw the optimization trajectory into the same plot and encode proximity to the plane as line thickness to determine when the optimizer moves away from it, as it is crucial for judging the degree of correspondence between landscape and trajectory.

For a quick overview of the optimization's descend speed, we employ a view that we call progression speed plot (view E of Fig. 1). This plot uses the metaphor of walking downhill (descending) to the minimum, and shows the remaining distance to travel for each optimization step (where the complete travel distance is normalized to 1). Task T2 is supported with this view.

Exploration of an optimization is enabled by interaction through selecting or cycling through optimization steps, constraints, or configurations, as well as through zooming and panning in the different views that respond to the selections made.

To select an optimization step, the user can either click into the progression speed and constraint plots, use a slider (on top of the GUI), or select a corresponding entry from an optimization log entry list (H in Fig. 1). The log entries give an overview of the inner workings of the optimization algorithm throughout the process, such as Newton steps taken, line search condition checked, or dual parameter updates done. We highlight line search backtracking steps in pink, graph query entries correspond to optimization steps. The naming of these entries stems from the structure of the log, output by the optimizer. Selecting individual joint configurations (robot time instants) can be done from list G in Fig. 1, which results in highlighting the corresponding trajectory in the path evolution view (A).

#### 4.2 Implementation

We implemented our visual analytics system as a desktop application in the Java programming language based on the AWT/Swing GUI environment. Libraries employed include Smile, EJML, JPlotter, jackson, and OkHttp. The solver runs in a separate software (KOMO [40]), producing optimization log files that our implementation facilitates to explore interactively. To obtain samples of the optimization space, we implemented an HTTP-server based interface in the motion planning framework, that runs simultaneously. We chose this method of data transfer due to its versatility and sustainability for future projects.

#### **CASE STUDY** 5

To showcase the usefulness of our system we present a case study and insights we could find with respect to a particular motion planning problem. We analyze an optimization run of a typical motion planning problem concerned with information need Q1. The robot in this problem is supposed to pick up and throw a ball so that it bounces of the ground and a wall to finally hit a target area.

When taking a look at the corresponding optimization in our visual analytics system, we would like to first get an overview. We typically start with checking the progression speed plot from which we can see the, in this case, large number of steps taken, which is around 1000 (Fig. 4). Examining the process' progression behavior (T2), we can observe that the optimizer is leaping forward

#### VINCI 2020, December 8-10, 2020, Eindhoven, Netherlands

#### Hägele et al.



Figure 4: Views examined in the first part of our case study. We observe slow progress throughout the majority of the optimization (a). We find an equality constraint group that increases over the process (b), and on expansion, see that several constraints of the group behave undesirably (c), i.e., not converging to zero.

in the first few iterations and then harshly slows down. Clicking the plot where we identified the sudden decrease in speed selects the corresponding optimization step (step 57 in this case). From there on, the optimizer crawls to its final position also hardly accelerating over the remaining iterations. This behavior is quite suspicious to us since this seems extremely ineffective, so we want to know what is going on.

Next, we take a look at the evolution of constraint values to check for any anomalies there (T3). Zooming in a little on the equality constraint line chart reveals a constraint that is increasing instead of decreasing during the slow part of the optimization (blue line in Fig. 4). The graph shows the maximum absolute value for the constraints sharing the same name, so to identify which constraints exactly are behaving in this way we expand this group of constraints to see the individuals (Fig. 4). From the expanded constraints, we see that most of them actually converge to zero as desired, but some of them rise up in the end and one is located far away from zero.

From the corresponding entries in the constraint tree, we can read off the robot time instants, i.e., the joint configurations to which the strange behaving constraints relate. Three consecutive configurations in the second part of the motion are related. This tells us that it is the part where the robot grabs the ball and throws it. Selecting the entry highlights the trajectories in the path evolution view. We use this to get an idea of the region on the path that is affected by the constraint (Fig. 5). Examining this part, we recognize a discontinuity in the motion path. This is due to the semantics of dimensions changing from one configuration to the next, indicating an event on the motion path that requires a different mathematical modeling. It tells us that the constraint is related to the exact moment of releasing the ball from its grip. This information could be valuable to the author of the motion problem when wanting to reformulate the problem to achieve better convergence.

We now want to go back to the original issue of slow progression. Checking the log table for frequent line searches that could cause slow progression reveals that there are no line searches taking place



Figure 5: Examining the part on the motion path that the misbehaving constraints relate to. We see a discontinuity in the path that corresponds to a change in configuration vector semantics. The constraints are related to the moment when the robot releases the ball.

in the largest portion of the process. As a next step, we consult the landscape view to take a look at the optimization trajectory (Fig. 6). We can observe that the optimizer actually moves away from a relatively better location in terms of its objective function indicated by the colored contours (yellow is less costly than green). This is not unusual in constraint optimization since the optimizer has to make a compromise to satisfy the constraints. Our hypothesis is that the suspected constraint forces the optimizer to leave this "cozy" place.

To test this hypothesis, we enable the display of constraint boundaries in this view (T4). We also make sure that the projection plane is representative of the slowly progressing part on the trajectory by

#### Visualization of Nonlinear Programming for Robot Motion Planning



Figure 6: Optimization landscape view examined in the last part of our case study. We zoom-in on the part of the optimization trajectory corresponding to slow progression we identified earlier (left). We can observe how the optimizer is climbing uphill in terms of the objective function (middle). When displaying the constraint boundaries, we see that the trajectory is moving in a parabola-shaped valley defined by the constraints (right).

selecting three points at steps 100, 600, and the final solution (Fig. 6). This orientates the plane to coincide with these three points so that we get a reasonable plane for our landscape. We can observe from this view that the trajectory moves in a parabolic shape and is actually enclosed by boundaries that have a similar shape. The areas shaded in gray correspond to infeasible regions of the suspected constraints for a certain feasibility threshold that we increased from one to six to be able to see a region at this large scale. Our interpretation of this is that the optimizer is trying to walk as closely as possible to the hyper surfaces defined by the constraint equations. The reason it keeps on walking is that it has not yet reached a location where it is close enough to all surfaces to be considered a feasible solution.

Considering the progression speed again and also no line search behavior for the majority of the trajectory, taking so many optimization steps still seems unreasonable. A sharper adaptive increase in step size was tested afterwards and led to a shortened time to convergence.

Since a single analysis scenario can only show some of the various issues that can occur in optimizations, we like to point the reader to our supplemental video for further impressions. In the video, we provide two other analysis scenarios while also demonstrating the system's interactive features, which are beyond the scope of this paper <sup>1</sup>.

## 6 DISCUSSION

In this section, we reflect on the project and report on the lessons learned during our research. NLP for motion planning is quite complex and comes with many quantities to visualize, such as objective costs, constraints, dual variables, gradient forces, and hypersurfaces – just to name a few. On top of that, it comes in highdimensional flavors and two notions of time (optimization time and robot time). We chose to focus on the evolutionary aspect, in terms of optimization time, to be able to follow the algorithm along its process. This choice turned out to work well for the assessment of long lasting optimizations and causes for that (Q1). The use of line charts to examine convergence behavior of constraints or loss is a widely used practice in optimization. They serve well as an entry point to optimization assessment, but experts want to be able to reason beyond the scope of such charts in order to understand what is actually happening. Leveraging linked views to allow for a more thorough exploration of the process worked well in our case.

During the course of this design study, we noticed that slow progression can be connected to ill conditioning and that long lasting optimizations often fail to produce feasible results. To assess optimizations with infeasible solutions (Q2), we also wanted to be able to reason about the behavior of the algorithm, analyze why it goes which way, and ideally see when it takes a "wrong" turn. Through the loss landscape visualization, we were able to provide means to achieve this. However, this technique is the most effective if the optimizer only moves in a low-dimensional subspace of the complete optimization space. This is not always the case and plane orientation becomes the deciding factor to effectively use this view.

To understand abstract high-dimensional points of optimization space, we chose a generic solution to display such time-involving intermediate solutions (motion paths in our case) by means of time curves. This works to some degree, however, experts would value a more concrete representation for their domain specific application, e.g., an animation of the robot performing the motion to better connect the abstract and physical world. We expect that integrating the animation capabilities of the motion planning framework into our system will greatly improve optimization space exploration in future work.

## 7 CONCLUSION

We presented a visual analytics system for the assessment of optimization processes of nonlinear constraint problems for robot motion planning. Except for the robot path evolution view, displaying the evolution of the motion path, the system provides domainagnostic views for the analysis of nonlinear programs. Leveraging a loss landscape visualization technique, the system allows users

<sup>&</sup>lt;sup>1</sup>https://doi.org/10.18419/darus-1128

VINCI 2020, December 8-10, 2020, Eindhoven, Netherlands

to have a look at the optimization trajectory and surrounding constraints even for the very high-dimensional problems we are facing in the domain. However, we found that the choice of projection plane is crucial for this technique to be effective. Better plane orientation could be explored in future research. The system's capabilities were showcased in a case study where we analyzed optimization runs of our collaborators' NLPs. Together with domain experts, we could confirm that we are able to gain new insights into optimizer behavior, detect flaws in the optimization process, and come up with issue resolving strategies by exploration and analysis with our system. Since most views are applicable to any nonlinear program, we think that our approach to nonlinear constraint optimization visualization will generalize to other problems. In future work, we plan to investigate respective applications and evaluate our approach in more depth.

## ACKNOWLEDGMENTS

The research has been supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2120/1 – 390831618, and the DAAD.

## REFERENCES

- W. Aigner, S. Miksch, H. Schumann, and C. Tominski. 2011. Visualization of Time-Oriented Data. Springer Science & Business Media.
- [2] David F Andrews. 1972. Plots of high-dimensional data. Biometrics 18, 1 (1972), 125–136.
- [3] G.S. Androulakis and M.N. Vrahatis. 1996. OPTAC: a portable software package for analyzing and comparing optimization methods by visualization. J. Comput. Appl. Math. 72, 1 (1996), 41–62.
- [4] B. Bach, P. Dragicevic, D. Archambault, C. Hurter, and S. Carpendale. 2014. A Review of Temporal Data Visualizations Based on Space-Time Cube Operations. In *EuroVis – STARs*. 23–41.
- [5] B. Bach, N. Henry-Riche, T. Dwyer, T. Madhyastha, J-D Fekete, and T Grabowski. 2015. Small MultiPiles: Piling time to explore temporal patterns in dynamic networks. *Computer Graphics Forum* 34, 3 (2015), 31–40.
- [6] B. Bach, C. Shi, N. Heulot, T. Madhyastha, T. Grabowski, and P. Dragicevic. 2016. Time Curves: Folding Time to Visualize Patterns of Temporal Evolution in Data. IEEE Transactions on Visualization and Computer Graphics 22, 1 (2016), 559–568.
- [7] Manuel Carro and Manuel Hermenegildo. 2000. Tools for constraint visualisation: The VIFID/TRIFID tool. In Analysis and Visualization Tools for Constraint Programming. Springer, 253–272.
- [8] Manuel Carro and Manuel Hermenegildo. 2000. Tools for search-tree visualisation: The apt tool. In Analysis and Visualization Tools for Constraint Programming. Springer, 237–252.
- [9] J. P. Charalambos and E. Izquierdo. 2001. Linear programming concept visualization. In Proceedings Fifth International Conference on Information Visualisation. 529–535.
- [10] A Chatterjee, P.P Das, and S Bhattacharya. 1993. Visualization in linear programming using parallel coordinates. *Pattern Recognition* 26, 11 (1993), 1725 – 1736.
- [11] P. C. Chen and Y. K. Hwang. 1998. SANDROS: a dynamic graph search algorithm for motion planning. *IEEE Transactions on Robotics and Automation* 14, 3 (1998), 390–403.
- [12] Herman Chernoff. 1973. The use of faces to represent points in k-dimensional space graphically. J. Amer. Statist. Assoc. 68, 342 (1973), 361–368.
- [13] John W Chinneck. 2006. Practical optimization: a gentle introduction. Systems and Computer Engineering, Carleton University, Ottawa (2006).
- [14] M. Ghoniem, H. Cambazard, J-D Fekete, and N. Jussien. 2005. Peeking in Solver Strategies Using Explanations Visualization of Dynamic Graphs for Constraint Programming. In Proceedings of the 2005 ACM Symposium on Software Visualization. 27–36.
- [15] M. Ghoniem, N. Jussien, and J-D Fekete. 2004. Visexp: visualizing constraint solver dynamics using explanations. In FLAIRS'04: Seventeenth international Florida Artificial Intelligence Research Society conference. 263–268.
- [16] Ian J. Goodfellow and Oriol Vinyals. 2015. Qualitatively characterizing neural network optimization problems. In 3rd International Conference on Learning Representations, ICLR.
- [17] S. Goodwin, C. Mears, T. Dwyer, M. G. de la Banda, G. Tack, and M. Wallace. 2017. What do Constraint Programming Users Want to See? Exploring the Role of

Hägele et al.

Visualisation in Profiling of Models and Search. *IEEE Transactions on Visualization* and Computer Graphics 23, 1 (2017), 281–290.

- [18] Henning Gruendl, Patrick Riehmann, Yves Pausch, and Bernd Froehlich. 2016. Time-Series Plots Integrated in Parallel-Coordinates Displays. *Computer Graphics Forum* 35, 3 (2016), 321–330.
- [19] S Heipcke. 1999. Comparing constraint programming and mathematical programming approaches to discrete optimisation - the change problem. *Journal of the Operational Research Society* 50, 6 (1999), 581–595.
- [20] Alfred Inselberg. 1985. The plane with parallel coordinates. The Visual Computer 1, 2 (1985), 69–91.
- [21] D. Jäckle, F. Fischer, T. Schreck, and D. A. Keim. 2016. Temporal MDS Plots for Analysis of Multivariate Data. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 141–150.
- [22] Eser Kandogan. 2000. Star coordinates: A multi-dimensional visualization technique with uniform treatment of dimensions. In Proceedings of the IEEE Information Visualization Symposium.
- [23] L. Kavraki and J. Latombe. 1994. Randomized preprocessing of configuration for fast path planning. In Proceedings of the IEEE International Conference on Robotics and Automation. 2138–2145, vol. 3.
- [24] L. E. Kavraki, P. Svestka, J. Latombe, and M. H. Overmars. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation* 12, 4 (1996), 566–580.
- [25] K Kondo. 1991. Motion planning with six degrees of freedom by multistrategic bidirectional heuristic free-space enumeration. *IEEE Transactions on Robotics and Automation* 7, 3 (1991), 267–277.
- [26] Harold W. Kuhn and Albert W. Tucker. 2014. Nonlinear Programming. Springer Basel, Basel, 247–258.
- [27] Jean-Claude Latombe. 2012. Robot Motion Planning. Vol. 124. Springer Science & Business Media.
- [28] Steven M LaValle. 2006. Planning Algorithms. Cambridge University Press.
- [29] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. 2018. Visualizing the Loss Landscape of Neural Nets. In Advances in Neural Information Processing Systems 31. 6389–6399.
- [30] S. Liu, D. Maljovec, B. Wang, P. Bremer, and V. Pascucci. 2017. Visualizing High-Dimensional Data: Advances in the Past Decade. *IEEE Transactions on Visualization and Computer Graphics* 23, 3 (2017), 1249–1268.
- [31] A. Messac and X. Chen. 2000. Visualizing the optimization process in real-time using physical programming. *Engineering Optimization* 32, 6 (2000), 721–747.
  [32] M. Noirhomme-Fraiture. 2002. Visualization of large data sets: the zoom star
- [32] M. Noirhomme-Fraiture. 2002. Visualization of large data sets: the zoom star solution. International Electronic Journal of Symbolic Data Analysis (2002), 26–39.
- [33] L. Nonato and A. Aupetit. 2018. Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE Transactions on Visualization and Computer Graphics* 25, 8 (2018), 2650–2673.
- [34] P. Pu and D. Lalanne. 2000. Interactive problem solving via algorithm visualization. In IEEE Symposium on Information Visualization. 145–153.
- [35] M. Sedlmair, M. Meyer, and T. Munzner. 2012. Design Study Methodology: Reflections from the Trenches and the Stacks. *IEEE Transactions on Visualization* and Computer Graphics 18, 12 (2012), 2431–2440.
- [36] Maxim Shishmarev, Christopher Mears, Guido Tack, and Maria Garcia De La Banda. 2016. Visual search tree profiling. *Constraints* 21, 1 (2016), 77–94.
- [37] Helmut Simonis, Paul Davern, Jacob Feldman, Deepak Mehta, Luis Quesada, and Mats Carlsson. 2010. A Generic Visualization Platform for CP. In International Conference on Principles and Practice of Constraint Programming. 460–474.
- [38] C. Tominski, J. Abello, and H. Schumann. 2004. Axes-based visualizations with radial layouts. In Proceedings of the 2004 ACM Symposium on Applied Computing. 1242–1247.
- [39] T. Torsney-Weir, T. Möller, M. Sedlmair, and R. M. Kirby. 2018. Hypersliceplorer: Interactive visualization of shapes in multiple dimensions. *Computer Graphics Forum* 37, 3 (2018), 229–240.
- [40] Marc Toussaint. 2014. Newton methods for k-order Markov Constrained Motion Problems. CoRR abs/1407.0414 (2014).
- [41] Marc Toussaint. 2015. Logic-Geometric Programming: An Optimization-Based Approach to Combined Task and Motion Planning. In Proceedings of the 24th International Conference on Artificial Intelligence. AAAI Press, 1930–1936.
- [42] Tea Tušar and Bogdan Filipič. 2014. Visualization of Pareto front approximations in evolutionary multiobjective optimization: A critical review and the prosection method. *IEEE Transactions on Evolutionary Computation* 19, 2 (2014), 225–245.
- [43] S. van den Elzen, D. Holten, J. Blaas, and J. J. van Wijk. 2016. Reducing Snapshots to Points: A Visual Analytics Approach to Dynamic Network Exploration. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 1–10.
- [44] Rainer Wegenkittl, Helwig Löffelmann, and Eduard Gröller. 1997. Visualizing the behaviour of higher dimensional dynamical systems. In Proceedings of the IEEE Conference on Visualization. 119–125.
- [45] Philip. Wolfe. 1969. Convergence Conditions for Ascent Methods. SIAM Rev. 11, 2 (1969), 226–235.
- [46] Pak Chung Wong and R Daniel Bergeron. 1994. 30 years of multidimensional multivariate visualization. *Scientific Visualization* 2 (1994), 3–33.