Quantifying Trajectory Safety by the Minimum Data Corruption Needed to Crash

Jared Miller[1] and Mario Sznaier[2]

[1] Automatic Control Laboratory (IfA), Department of Information Technology and Electrical Engineering (D-ITET), ETH Zürich, Physikstrasse 3, 8092, Zürich, Switzerland,
jarmiller@control.ee.ethz.ch
[2] Robust Systems Lab, ECE Department, Northeastern University, Boston, MA 02115, msznaier@coe.neu.edu

Abstract. This work quantifies the safety of trajectories of a dynamical system by the perturbation intensity required to render a system unsafe (crash into the unsafe set). Computation of this measure of safety is posed as a peak-minimizing optimal control problem. Convergent lower bounds on the minimal peak value of controller effort are computed using polynomial optimization and the moment-Sum-of-Squares hierarchy. The crash-safety framework is extended towards data-driven safety analysis by measuring safety as the maximum amount of data corruption required to crash into the unsafe set.

Keywords: Safety, Peak, Optimal Control, Moment-Sum-of-Squares, Data-Driven

1 Introduction

A trajectory starting at an initial point $x_0 \in X$ following dynamics $\dot{x} = f_0(t, x)$ is safe with respect to the unsafe set X_u in the time horizon $t \in [0, T] \subset [0, \infty)$ if there does not exist a time t' such that $x(t' \mid x_0)$ is a member of X_u. The set X_0 is safe with respect to X_u if all initial points $x_0 \in X_0$ generate safe trajectories. This abstract quantifies the safety of trajectories by maximum control effort (Optimal Control Problem (OCP) cost) needed to crash the agent into the unsafe set. An example of this type of safety result is if tilting a car's steering wheel by a maximum extent of 3° over the course of its motion would cause the car to crash. In a data-driven framework, a continuous-time trajectory is labeled safe if it would require the true system to have a large constraint violation against any of its state-derivative data observations in order to crash. The process of analyzing safety by peak-minimizing-OCP cost will be referred to as 'crash safety'.

1.1 Prior Work

Peak-minimizing control problems are a particular form of robust optimal control in which the minimizing agents are $(t, x_0, w(\cdot))$ and the maximizing agent is

$t' \in [0,t]$. Necessary conditions for these robust programs may be found in (21). Instances of peak-minimizing control include minimizing the maximum number of infected persons in an epidemic under budget constraints (17) and choosing flight parameters to minimize the maximum skin temperature during atmospheric reentry (12; 7).

This paper continues a sequence of research about quantifying the safety of trajectories. Unsafety can be proven using path-planning by finding a feasible pair $(t', x_0) \in [0,T] \times X_0$ such that $x(t' \mid x_0) \in X_0$. Barrier (19; 18) and Density functions (20) are binary certificates confirming that there does not exist an unsafe trajectory based on the satisfaction of nonnegativity constraints. Safety margins use maximin peak estimation to estimate the X_u-representing-inequality-constraint violation (13). The distance of closest approach between a trajectory starting in X_0 and points in X_u is a more interpretable measure of safety than abstract safety margins (14). Even so, distance estimation does not tell the full story; a trajectory may lie close to X_u in the sense of distance, but it could require a large value of Q^* to render the same trajectory unsafe.

Direct solution of OCPs using the Hamilton-Jacobi-Bellman (HJB) equation or the Pontryagin Maximum Principle may be challenging, especially when solutions do not exist in closed form (10). These generically non-convex OCPs may be lifted into convex infinite-dimensional Linear Programs (LPs) in occupation measures (9), whose dual LP involve subvalue functions satisfying HJB inequalities. These infinite-dimensional LPs produce lower-bounds on the true OCP, with equality holding under compactness and regularity conditions. The moment-Sum of Squares (SOS) hierarchy of Semidefinite Programs (SDPs) may be used to produce a rising sequence of lower bounds to the true OCP if all problem aspects (cost, dynamics, sets) are polynomial-representable Basic Semi-algebraic (BSA) (5). This infinite-dimensional LP and finite-dimensional SDP pattern has also been applied to reachable set estimation (4), peak estimation (3), and maximum controlled invariant set estimation (6).

The extended work in (16) includes proofs of convergence, proofs of strong duality, discussion about subvalue functions to map data corruption, and additional examples.

1.2 Notation

The set of real numbers is $\mathbb{R}$ and the n-dimensional real vector spaces is $\mathbb{R}^n$. The all-ones vector is $\mathbf{1}$. The set of natural numbers is $\mathbb{N}$ and the set of n-dimensional multi-indices is $\mathbb{N}^n$. The set of natural numbers between a and b is $a..b \subset \mathbb{N}$. The cone of $n \times n$ symmetric Positive Semidefinite (PSD) matrices is $\mathbb{S}^n_+$.

The set of polynomials of an indeterminate x with real-valued coefficients is $\mathbb{R}[x]$. The degree of a polynomial $p \in \mathbb{R}[x]$ is $\deg p$. The vector space of polynomials up to degree $d \in \mathbb{N}$ is $\mathbb{R}[x]_{\leq d}$. The coefficients of a polynomial $p \in \mathbb{R}[x]$ are $\mathrm{coeff}_x(p(x))$.

The ring of continuous functions over a space $S \subseteq \mathbb{R}^n$ is $C(S)$. The set of first-differentiable functions over S is $C^1(S) \subset C(S)$. The subcone of nonnegative functions over S is $C_+(S) \subset C(S)$.

2 Data-Driven Crash-Safety Analysis

This section motivates crash-safety in the context of data-driven analysis.

2.1 Data-Driven Overview

In this section, we will assume that N_s time-state-derivative data records $\mathcal{D} = \{(t_k, x_k, y_k)\}_{k=1}^{N_s}$ are provided for the true system $\dot{x} = F(t, x)$. The data records in $\mathcal{D}$ are corrupted by L_∞-bounded noise of intensity ϵ with

$$\forall k = 1..N_s \qquad\qquad \|y_k - F(t_k, x_k)\|_\infty \leq \epsilon. \tag{1}$$

We are given a dictionary of functions $(f_0, \{f_\ell\}_{\ell=1}^L)$ that are Lipschitz in $[0, T] \times X$ (e.g. monomials). We are also given the knowledge that there exists at least one ground-truth choice of parameters $w^* \in \mathbb{R}^L$ with

$$F(t, x) = f_0(t, x) + \textstyle\sum_{\ell=1}^L w_\ell^* f_\ell(t, x). \tag{2}$$

In the L_∞-bounded polytopic framework, the crash-safety problem finds an infimal upper bound on the data corruption needed to crash into the unsafe set:

$$\begin{aligned}
Z^* = \inf_{t,\, x_0,\, z,\, w} \; & z && \text{(3a)}\\
& \forall t' \in [0, T] : \;\; \dot{x}(t') = f_0(t', x) + \textstyle\sum_{\ell=1}^L w_\ell f_\ell(t', x(t')), \quad \dot{z}(t') = 0 && \text{(3b)}\\
& x_0 \in X_0, \;\; x(t \mid x_0, w) \in X_u && \text{(3c)}\\
& \forall k = 1..N_s : \quad z \geq \|f_0(t_k, x_k) + \textstyle\sum_{\ell=1}^L w_\ell f_\ell(t_k, x_k) - y_k\|_\infty && \text{(3d)}\\
& z \in Z, \;\; w \in \mathbb{R}^L, \;\; t \in [0, T]. && \text{(3e)}
\end{aligned}$$

If the returned value of (3) is $Z^* = 0$, then there exists some choice of model parameters w that exactly fit the data $\mathcal{D}$ by (2). Additionally, this choice w renders at least one trajectory $x(\cdot)$ starting from X_0 is unsafe (crashes into X_0). Values of Z^* greater than 0 are a certificate of safety in the model structure. A larger value of Z^* indicates that the data must be increasingly corrupted in order to render any trajectory unsafe. Safety is certified if $Z^* > \epsilon$, though we note that the true value of ϵ may be a-priori unknown.

2.2 Data Representation

For each $k = 1..N_s$, define the data-record matrices $\Gamma_k, \;\; h_k$ by

$$\Gamma_k = \big[f_1(t_k, x_k), \cdots, f_L(t_k, x_k)\big] \qquad\qquad h_k = f_0(t_k, x_k) - y_k. \tag{4}$$

Letting Γ and h be the vertical concatenations of $\{\Gamma_k\}$ and $\{h_k\}$ respectively, we can define the L_∞ performance function and support set as

$$J(w) = \|\Gamma w - h\|_\infty, \quad Z = [0, J_{\max}], \tag{5}$$

and the support set for (w, z) from (3d) as

$$\Omega = \left\{(w, z) \in \mathbb{R}^L \times Z : \begin{matrix} \Gamma w \leq z\mathbf{1} - h \\ -\Gamma w \leq z\mathbf{1} + h \end{matrix}\right\}. \tag{6}$$

2.3 Standard Crash-Safety Linear Program

This section will convert the generically nonconvex optimal control problem in (3) into an infinite-dimensional LP in continuous functions using the methods of (9). Let $\mathcal{L}_f$ be the Lie derivative associated with f for $v(t,x,z) \in C^1$ as

$$\mathcal{L}_f v(t,x,z,w) = (\partial_t + f_0(t,x)\cdot\nabla_x)v(t,x,z) + \sum_{\ell=1}^{L} w_\ell f_\ell(t,x)\cdot\nabla_x v(t,x,z) \tag{7}$$

An auxiliary function $v \in C^1$ may be defined to form an LP formulation of the crash-safety OCP in (3):

$$\begin{aligned}
q^* = \sup_{\gamma\in\mathbb{R},\ v}\ & \gamma && && (8a)\\
& v(0,x,z) \geq \gamma && \forall (x,z) \in X_0 \times Z && (8b)\\
& v(t,x,z) \leq z && \forall (t,x,z) \in [0,T]\times X_u \times Z && (8c)\\
& \mathcal{L}_f v(t,x,z,w) \geq 0 && \forall (t,x,z,w) \in [0,T]\times X \times \Omega && (8d)\\
& v(t,x,z) \in C^1([0,T]\times X \times Z). && && (8e)
\end{aligned}$$

Theorem 1 *Under assumptions A1-A5, programs* (3) *and* (8) *will have equal objectives* $q^* = Q^*$.

Proof. Program (3) with optimum Z^* is a standard-form OCP with free terminal time and zero running cost. Under assumptions A1-A5, Theorem 2.1 of (9) proves that $Z^* = q^*$. Section 6.3 of (9) specifically discusses state-dependent controls (e.g. $(w,z) \in \Omega$).

The infinite-dimensional program in (8) must be discretized into a finite-dimensional convex optimization problem in order to admit computation. One such method to perform this discretization is through the moment-SOS hierarchy, in which $v(t,x,z)$ is restricted to be a polynomial, and the inequality constraints (8b)-(8d) are replaced by Putinar Positivestellensätze (8). Assuming that f_0 and each f_ℓ are polynomials in (t,x), define $\tilde{d} = d + \max_{\ell\in 0..L} \lfloor \deg f_\ell/2 \rfloor$ as the dynamics degree of (2) given a degree $d \in \mathbb{N}$. Imposition of constraint (8d) given a polynomial v of degree 2d using the moment-SOS hierarchy requires a maximal-size PSD matrix constraint of dimension $\binom{n+L+2+\tilde{d}}{\tilde{d}}$, which is intractably large as L increases (more dictionary entries).

2.4 Robust Crash-Safety Linear Program

We will use the input-affine structure of dynamics and polytopic form of (3d) to form an LP that eliminates the uncertainty w. This elimination leads to increasingly tractable SOS-based SDPs with a maximal PSD matrix size $\binom{n+2+\tilde{d}}{\tilde{d}}$, because the maximal-size PSD matrices will no longer depend on L.

We will eliminate the w variable from (8d) by introducing new nonnegative multiplier functions $\{\zeta^+, \zeta_j^-\}_{j=1}^{2nT}$. This elimination proceeds using the infinite-dimensional robust counterpart method of (15), which requires that (8d) hold strictly (with a '> 0') constraint.

Theorem 2 *A strict version of Lie constraint in* (8d) *may be robustified (will have the same feasibility/infeasibility conditions) into*

$$\forall (t,x,z) \in [0,T] \times X \times Z : \\ \mathcal{L}_{f_0} v - (z\mathbf{1} - h)^T \zeta^+ - (z\mathbf{1} + h)^T \zeta^- > 0 \tag{9a}$$

$$\forall \ell = 1..L : \quad (\Gamma^T)_\ell (\zeta^+ - \zeta^-) + f_\ell \cdot \nabla_x v = 0 \tag{9b}$$

$$\forall j = 1..2nT : \ \zeta_j^+, \zeta_j^- \in C_+([0,T] \times X \times Z). \tag{9c}$$

Proof. See the Proof of Theorem 4.1 of the extended version (16).

Remark 1 *Theorems 4.2 and 4.4 of (15) may be applied when Ω is a more general semidefinite representable set parameterized by z, such as an intersection of ellipsoids for L_2-bounded noise, or a projection of spectahedra for semidefinite bounded noise.*

3 Examples

This section demonstrates the utility of the crash-safety framework. Robust decompositions of the Lie constraint are applied in all examples. MATLAB R2021a code to generate examples is available at https://github.com/Jarmill/crash-safety. All SDP are generated using YALMIP (11) and solved using Mosek (2). Finite-degree crash-bounds SOS truncations of (8) with robust Lie constraint (9) are compared against OCP bounds found using the solver CasADi (1).

Data $\mathcal{D}$ is collected for the Flow system from (20):

$$\dot{x} = \begin{bmatrix} x_2 \\ -x_1 - x_2 + \frac{1}{3}x_1^3 \end{bmatrix}. \tag{10}$$

A total of $N_s = 40$ samples with perfect knowledge in dynamics $\dot{x}_1 = x_2$ and a ground-truth noise bound of $\epsilon = 0.5$ in the coordinate $\dot{x}_2$ are collected. The noisy derivative data in $\mathcal{D}$ and ground-truth derivatives are drawn in the orange and blue arrows respectively in Figure 1a. It is assumed that $\dot{x}_2$ is described by a cubic polynomial in (x_1, x_2). The parameterized polytope $\{w \mid Aw \leq b + z\}$ (Ω with fixed z value) has $L = 10$ dimensions and $m = 2nT = 80$. The minimum possible corruption while obeying (2) under the cubic noise model is $\inf_{(w,z)\in\Omega} z = 0.4617$.

The crash-safety problem (8) was solved with the unsafe set $X_u = \{x \mid (x_1 + 0.25)^2 + (x_2 + 0.7)^2 \leq 0.5^2,\ (0.95 + x_1 + x_2)/\sqrt{2} \leq 0\}$ between $t = [0, 5]$ time units in the space $X = \{x \in \mathbb{R}^2 : \|x\|_2^2 \leq 8\}$. Table 1 reports bounds for the crash-corruption $Q(X_0)$ by solving Lie-robustified SOS tightenings of (8) from degrees 1..4 with $J_{\max} = 1$

Safety of trajectories starting in X_0 is certified because the crash-bound $\tilde{q}_4^* = 0.5499$ is greater than the ground-truth noise-bound $\epsilon = 0.5$. Figure 1b uses the CasADi optimal control suite (1) to numerically solve the crash program (3). The numerical crash-bound of $q^{\text{CasADI}} = 0.5499$ is approximately equal (up to four decimal places) to the crash-bound $q_4^* = 0.5499$.

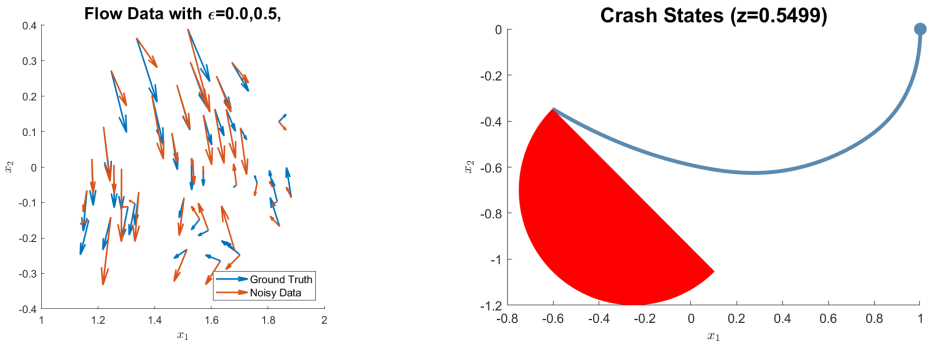

(a) Observed data (10) (b) CasADi- computed crash-bound

Fig. 1: Crash-Safety Analysis for the data-driven Flow system (10).

Table 1: Data-Driven Crash-bounds at $X_0 = [1; 0]$ under SOS tightenings

order	1	2	3	4
specific (8)	0.0582	0.4423	0.4864	0.5499

These crash-bounds should be compared against the L_2 distance estimates of $c^*_{1:5} = [1.698 \times 10^{-5},\ 0.1936,\ 0.2003,\ 0.2009,\ 0.2013]$ from Section 6.3 of (15). The distance estimates do not indicate that adding an additional budget of 0.0499 constraint violation will cause at least one trajectory to enter the unsafe set.

4 Conclusion

This paper utilized peak minimizing control in order to perform safety analysis. The returned values from SOS programs are lower-bounds on the minimum data corruption needed to crash into the unsafe set. Crash-safety adds a new perspective on the safety of trajectories, covering some of the blind spots of distance estimation and safety margins.

Future work involves attempting to reduce computational burden finite-dimensional truncations of the Crash programs by identifying new kinds of structure (in addition to robust decompositions) to hopefully allow for real-time computation. Other extensions could include applying these methods to other classes of systems (e.g., discrete-time, hybrid), and creating a stochastic interpretation of crash-safety.

Bibliography

[1] Andersson, J.A., Gillis, J., Horn, G., Rawlings, J.B., Diehl, M.: CasADi: a software framework for nonlinear optimization and optimal control. Mathematical Programming Computation **11**, 1–36 (2019)

[2] ApS, M.: The MOSEK optimization toolbox for MATLAB manual. Version 9.2. (2020). URL https://docs.mosek.com/9.2/toolbox/index.html

[3] Fantuzzi, G., Goluskin, D.: Bounding Extreme Events in Nonlinear Dynamics Using Convex Optimization. SIAM Journal on Applied Dynamical Systems **19**(3), 1823–1864 (2020)

[4] Henrion, D., Korda, M.: Convex Computation of the Region of Attraction of Polynomial Control Systems. IEEE TAC **59**(2), 297–312 (2013)

[5] Henrion, D., Lasserre, J.B., Savorgnan, C.: Nonlinear optimal control synthesis via occupation measures. In: 2008 47th IEEE Conference on Decision and Control, pp. 4749–4754. IEEE (2008)

[6] Korda, M., Henrion, D., Jones, C.N.: Convex computation of the maximum controlled invariant set for polynomial control systems. SIAM Journal on Control and Optimization **52**(5), 2944–2969 (2014)

[7] Kreim, H., Kugelmann, B., Pesch, H.J., Breitner, M.H.: Minimizing the Maximum Heating of a Reentering Space Shuttle: An Optimal Control Problem with Multiple Control Constraints. Optimal Control Applications and Methods **17**(1), 45–69 (1996)

[8] Lasserre, J.B.: Moments, Positive Polynomials And Their Applications. Imperial College Press Optimization Series. World Scientific Publishing Company (2009)

[9] Lewis, R., Vinter, R.: Relaxation of Optimal Control Problems to Equivalent Convex Programs. Journal of Mathematical Analysis and Applications **74**(2), 475–493 (1980)

[10] Liberzon, D.: Calculus of Variations and Optimal Control Theory: A Concise Introduction. Princeton university press (2011)

[11] Lofberg, J.: YALMIP : a toolbox for modeling and optimization in MATLAB. In: ICRA (IEEE Cat. No.04CH37508), pp. 284–289 (2004)

[12] Lu, P., Vinh, N.X.: Minimax optimal control for atmospheric fly-through trajectories. Journal of Optimization Theory and Applications **57**(1), 41–58 (1988)

[13] Miller, J., Henrion, D., Sznaier, M.: Peak Estimation Recovery and Safety Analysis. IEEE Control Systems Letters **5**(6), 1982–1987 (2020). DOI 10.1109/LCSYS.2020.3047591

[14] Miller, J., Sznaier, M.: Bounding the Distance to Unsafe Sets with Convex Optimization (2021). ArXiv: 2110.14047

[15] Miller, J., Sznaier, M.: Analysis and Control of Input-Affine Dynamical Systems using Infinite-Dimensional Robust Counterparts (2023). Arxiv:2112.14838

[16] Miller, J., Sznaier, M.: Quantifying the safety of trajectories using peak-minimizing control. arXiv preprint arXiv:2303.11896 (2023)
[17] Molina, E., Rapaport, A.: An optimal feedback control that minimizes the epidemic peak in the SIR model under a budget constraint. Automatica **146**, 110,596 (2022)
[18] Prajna, S.: Barrier certificates for nonlinear model validation. Automatica **42**(1), 117–126 (2006)
[19] Prajna, S., Jadbabaie, A.: Safety Verification of Hybrid Systems Using Barrier Certificates. In: International Workshop on Hybrid Systems: Computation and Control, pp. 477–492. Springer (2004)
[20] Rantzer, A., Prajna, S.: On Analysis and Synthesis of Safe Control Laws. In: 42nd Allerton Conference on Communication, Control, and Computing, pp. 1468–1476. University of Illinois (2004)
[21] Vinter, R.B.: Minimax optimal control. SIAM journal on control and optimization **44**(3), 939–968 (2005)