Hilbert scheme of points, Local uniformization and Reinforcement learning

Honglu Fan

University of Geneva

2022

It is not very common to see these three words together in a talk's title. I would like to give a short tour about how they come together in a joint project with Gergely Berczi (Aarhus University), as well as Mingcong Zeng (Max Plank Institute).

We are still seeking more applications and wanting to build stronger RL models...

Global singularity theory and singularity of maps

Our original motivation comes from a specific need in Berci's work on global singularity theory.

(Adapted from [Bé20]) Let $f : N \to M$ be a map between two complex manifolds who has $\dim(N) < \dim(M)$. $p \in N$ is a singular point of f if the rank of the differential $df_p : T_pN \to T_{f(p)}M$ is less than $\dim(N)$.
Locally near p, the map-germ $f_p : (\mathbb{C}^n, 0) \to (\mathbb{C}^m, 0)$ can be thought of as m power series with n variables, thus forming an infinite dimensional moduli space $J(n, m)$.
Change of coordinates induces group actions Diff_n on $(\mathbb{C}^n, 0)$, and Diff_m on $(\mathbb{C}^m, 0)$. Thus, $\mathrm{Diff}_n \times \mathrm{Diff}_m$ acts on $J(n, m)$.

Global singularity theory, cont.

- Moduli of map-germ: $J(n, m)$.
- Group action by: $\mathrm{Diff}_n \times \mathrm{Diff}_m$.
- Singularity type of maps should be something invariant under change of coordinates $\mathrm{Diff}_n \times \mathrm{Diff}_m$

Definition

A singularity $O \subset J(n, m)$ is a $\mathrm{Diff}_n \times \mathrm{Diff}_m$-invariant subset.

Definition

Given a holomorphic map $f : N \to M$ and a singularity O, the locus of singularity is $\Sigma_O(f) = \{p \in N; f_p \in O\}$.

Theorem (Thom(real),Damon(complex))

The homology class $[\Sigma_O(f)]$ for generic f is given by a bivariant class τ_O given by Chern characteristics of $TN, f^\star TM$.

Thom polynomials

In short, given a singularity $O \subset J(n, m)$, there is the universal polynomial called **Thom polynomial** in Chern characters of TN, TM which is independent of N, M and the map f.

A_k**-singularities (Morin singularities):**

$$\Theta_k^{n \to m} = \{f : (\mathbb{C}^n, 0) \to (\mathbb{C}^m, 0) | \mathbb{C}[x_1, \cdots, x_n]/(f_1 \cdots, f_m) \cong \mathbb{C}[t]/t^k\}$$

Question: What is the Thom polynomial of A_k singularities?

Berczi-Szenes and Berczi's approach in [BS12][Bé20]

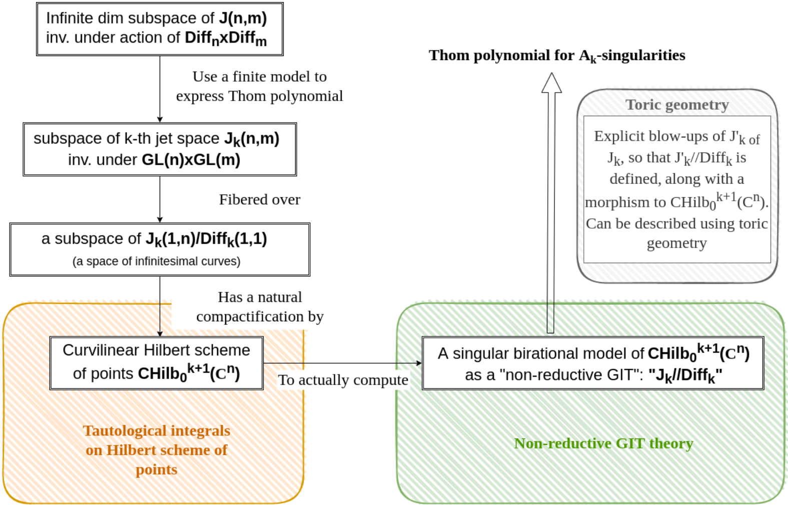

Punchline

Computing Thom classes for Morin singularities amounts to understanding specific sequences of toric blow-ups on a birational model of **curvilinear Hilbert scheme of points**.

Understanding the blow-up sequences in a resolution becomes the key. Local version of resolution of singularity is called **local uniformizations**.

Valuations

Let K be a field, Γ an Abelian group with total ordering. Let $K^\star$ denote the multiplicative group of K.

Definition

A valuation of K with value group Γ is a surjective group homomorphism $\nu : K^\star \to \Gamma$ such that for all $x, y \in K^\star$,

$$\nu(x + y) \geq min\{\nu(x), \nu(y)\}$$

Local Uniformization

Let X be an irreducible algebraic variety, $\xi \in X$ a point of X and ν a valuation of $K = K(X)$.

Lemma (Consequence of valuative criteria)

Let $\pi : X' \to X$ be a birational projective morphism, there exists a unique point $\xi' \in \pi^{-1}(\xi)$ such that $\nu(\mathcal{O}_{X',\xi'}) \geq 0$ and $\nu(\mathfrak{m}_{X',\xi'}) > 0$ (ν is centered at ξ').

Definition

A local uniformization of X with respect to ν is a birational projective morphism $\pi : X_0 \to X$ such that the center ξ_0 of ν in $\pi^{-1}(\xi)$ is a regular point of X_0.

The geometry

Local uniformizations are sometimes seen as a local version of resolution of singularity.
(In other words, for a local ring coming from an algebraic variety, there exists a regular local ring that is "birationally equivalent" to it)

Characteristic 0: proven by Zariski.
Characteristic p in dimension ≥ 4: open.

An important operation in resolution of singularities and local uniformizations is the blow-up.

Blow up

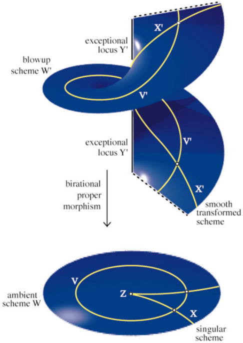

-The Hironaka Theorem on Resolution of Singularities (Or: A proof we always wanted to understand), H. Hauser

Now

Where is the machine learning part?

Local uniformization as a double-player game

(Hironaka's polyhedra game)
Let us play a game. There are two players. Let's call them **host** and **agent** (my own language, not in Hironaka's literature).
The game state is represented by a set of points $S \in \mathbb{Z}^n$ who are the vertices of the Newton polytope S itself. Every turn, host makes a move, agent makes a move, then the game state changes and we repeat the process.
Every turn,

1. The host chooses a subset $I \subset \{1, 2, \cdots, n\}$ such that $|I| \geq 2$.
2. The agent chooses a number $i \in I$.

Hironaka's polyhedra game, cont.

Combining $i \in I \subset [n]$, we change the 'state' S to the next step according to the following:

$$x_j \mapsto \begin{cases} x_j, & \text{if } i \neq j \\ \sum_{k \in I} x_k, & \text{if } i = j \end{cases},$$

for points $(x_1, \cdots, x_n) \in \mathbb{Z}^n$. We subsequently apply Newton polytope to the transformed points and only keep the vertices.

A state is terminal if it consists of one single point. In this case, the game will not continue and the host wins.
As a result,

1. the host wants to reduce the number of S and stop the game as quickly as possible,
2. the agent wants to keep the number of S for as long as possible (potentially forever).

Winning host as a local uniformization solver

Dictionary between the game and the math:

- The game state: a monomial ideal generated by monomials corresponding to points.
- Host choice: a toric blow-up.
- Agent choice: pick a chart on the exceptional locus.
- Changing game states: coordinate change of blow-up.
- A host who wins against any agent in finite steps: a local uniformization of the starting monimial ideal!
- A host who wins faster: a local uniformization that requires fewer toric blow-ups.

Theorem ([Spi83]; [Zei06])

Winning strategies for hosts exist, but are not unique.

Variations

"Hard" Hironaka's polyhedra game:

Could imply (not equivalent to!) local uniformization in any characteristics but counterexamples are found (for some games and some agents, hosts cannot win).

Weighted polyhedra game:

Models off the blow-up sequence in J_k.

Maybe there are more corresponding to other singularities...

The catch

Spivakovsky and Zeillinger's winning strategies sometimes requires too many steps to win, i.e., for some singularities, their methods take way too many blow-ups to resolve (including in the computation of Thom polynomials for Morin singularities). We need to go beyond them.

Now, what comes to your mind when you want to train superhuman game players?

AlphaGo, AlphaZero

Mastering the game of Go with Deep Neural Networks & Tree Search
David Silver, Aja Huang, Christopher Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, J Nham *, Nal Kalchbrenner, I Sutskever *, Thore Graepel, Timothy Lillicrap, Maddy Leach, Koray Kavukcuoglu, Demis Hassabis

Mastering the game of Go without Human Knowledge
David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, Demis Hassabis

A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play
David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, Demis Hassabis

Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model
David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, Demis Hassabis

HironakaZero???

Key algorithms: deep Q-learning, PPO, Monte Carlo Tree Search

The subject of reinforcement learning has a long history and dates back to Bellman's dynamic programming (1958). There are many many names that shaped the subject but the current look of this subject was formed when Chris Watkins's Q-learning (1989) unifies previous branches. Later, there is the popular policy gradient algorithm PPO. When sample efficiency is not a problem, Monte Carlo tree search is also very powerful and was the workhorse of AlphaZero.

We are working on RL training a local uniformization solver on GPU, which hopefully solves the computation of Thom polynomials for Morin singularities and more. https://arxiv.org/abs/2307.00252 is a summary of what we have so far.

Some resolutions

D5

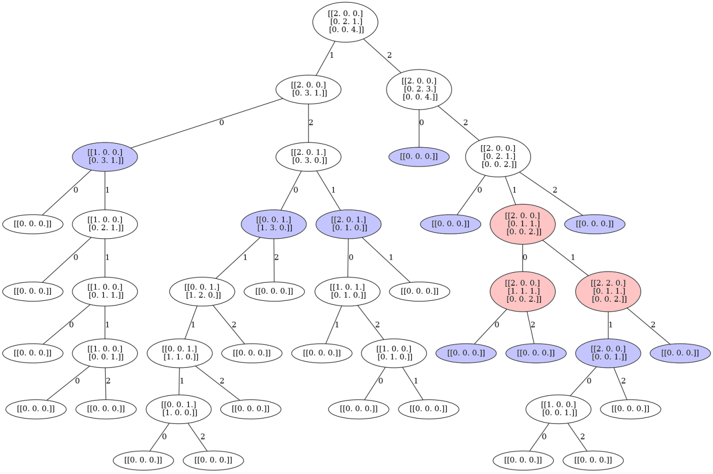

Some resolutions

E8

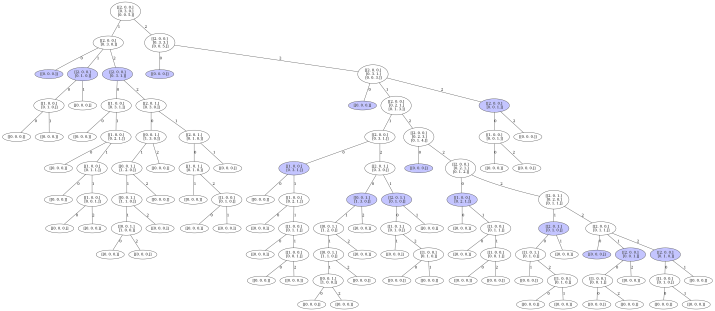

Special thank

Special thank to High-Flyer who generously shares their 10-million-dollar GPU cluster with us.

References

- Gergely Bérczi and András Szenes, *Thom polynomials of Morin singularities*, Ann. Math. (2) **175** (2012), no. 2, 567–629 (English).
- Gergely Bérczi, *Non-reductive geometric invariant theory and thom polynomials*.
- Mark Spivakovsky, *A solution to hironaka's polyhedra game*, Arithmetic and geometry, Springer, 1983, pp. 419–432.
- Dominik Zeillinger, *A short solution to hironaka's polyhedra game*, ENSEIGNEMENT MATHEMATIQUE **52** (2006), no. 1/2, 143.