Data-Driven Encodings for Robust, Scalable, and Interpretable Evolutionary Computation

Sebastian Risi

IT UNIVERSITY OF COPENHAGEN

modl.ai

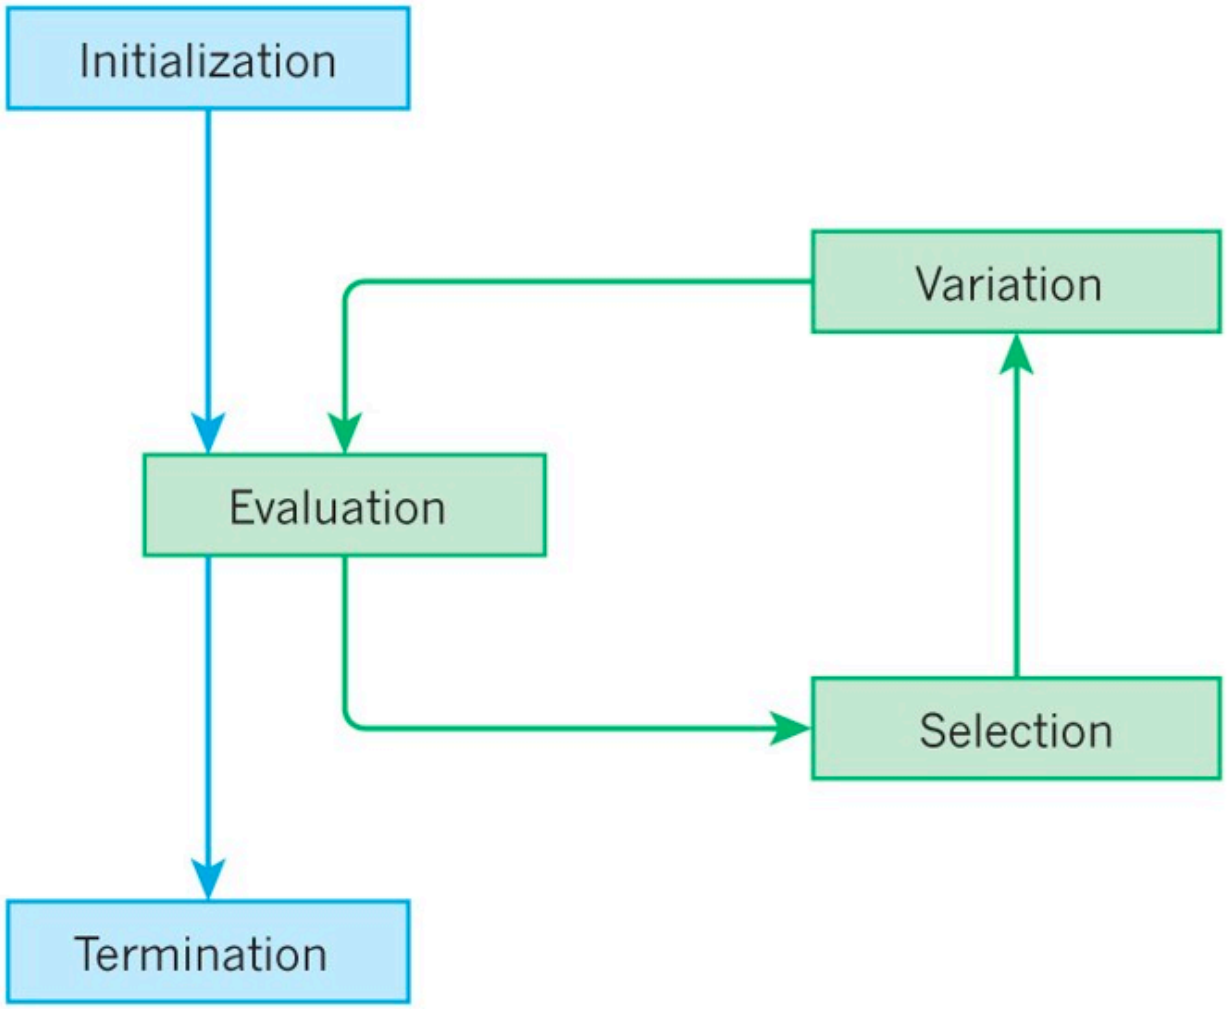

Eiben, A., Smith, J. From evolutionary computation to the evolution of things. *Nature* **521,** 476–482 (2015).

Direct vs. Generative Encodings

- Direct Encoding: Each genotypic element specifies an independent phenotypic element

- Generative Encoding: Genotypic elements can influence many phenotypic elements
  - Hand-designed encoding (e.g. CPPN, Hypernetworks)
  - Learn them from data? (the ML way)

Compositional Pattern Producing Networks (CPPNs; Stanley 2007)

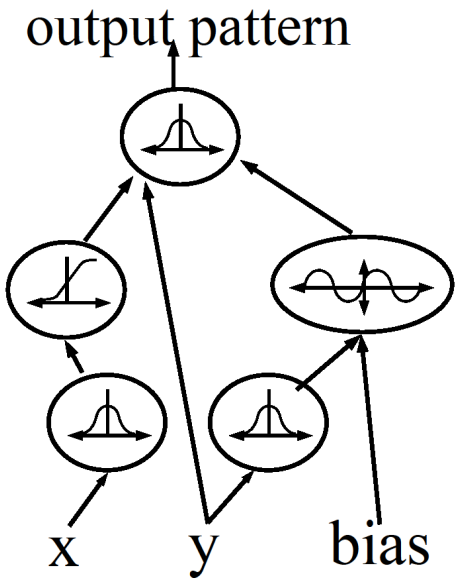

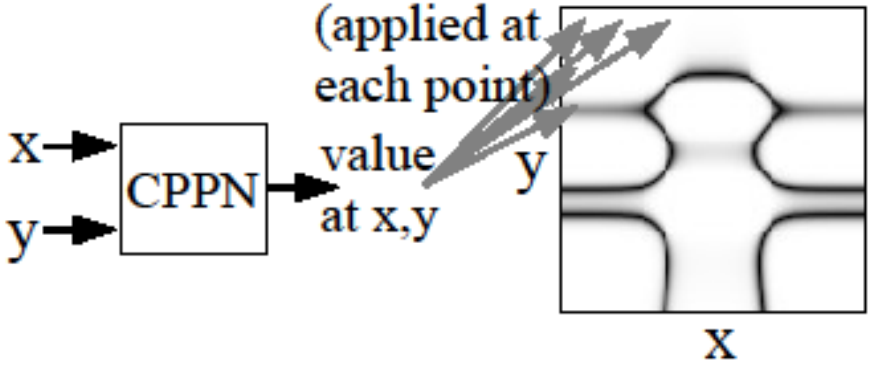

+ Strong bias towards regularities that we might want to see
- Hard to evolve towards a particular target (Woolley & Stanley, 2011)
- Need to be tailored to a specific domain

A Flower Painting CPPN

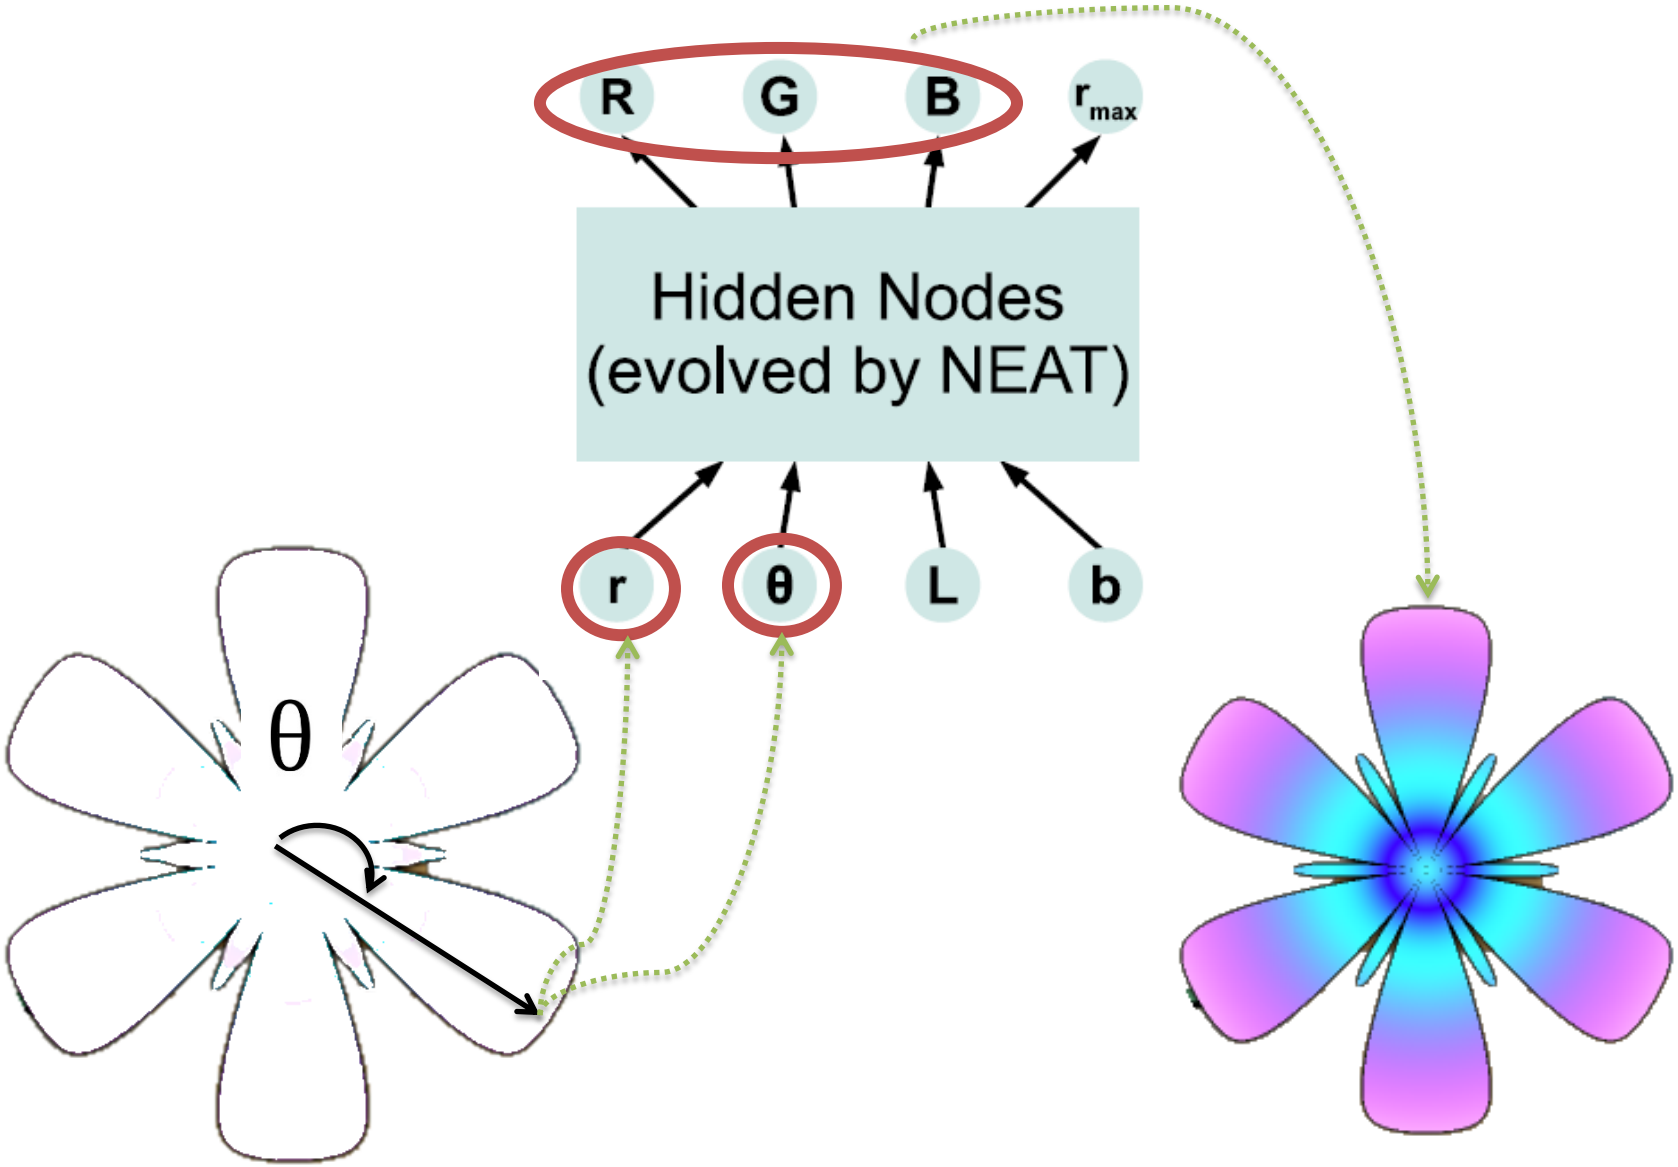

Flower Evolution: Pollinating a Flower

Planting the Offspring

Latent Variable Evolution (LVE)

- A learned compact genotype-to-phenotype mapping ➔ robust mutations
- Applicable to many different domains

Bontrager, Togelius, Memon 2017

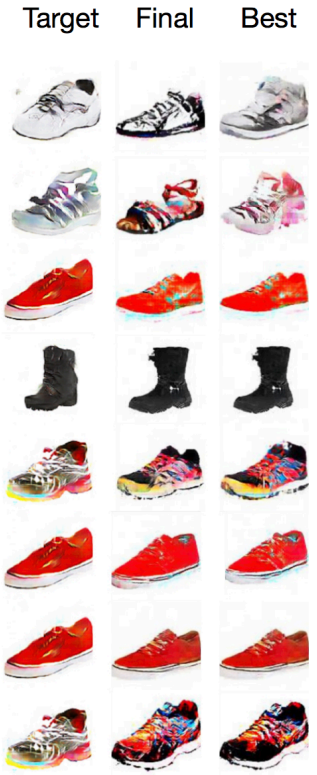

Bontrager, Lin, Togelius, Risi, 2018

Can be particularly useful if there already exists a large corpus of content we want to emulate

Generative and Adversarial Networks (GANs)

Goodfellow 2014

NVIDIA 2017

Radford et al. 2015

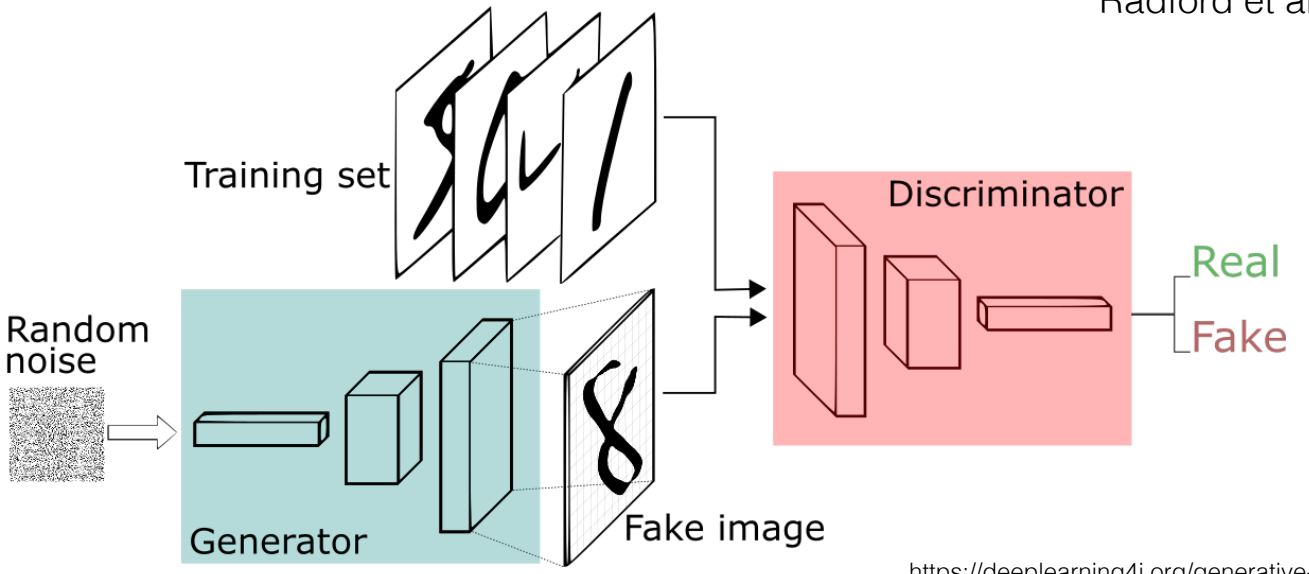

https://deeplearning4j.org/generative-adversarial-network

Evolving Mario Levels in the Latent Space of a Deep Convolutional Generative Adversarial Network

Volz, Schrum, Liu, Lucas, Smith, Risi, GECCO 2018

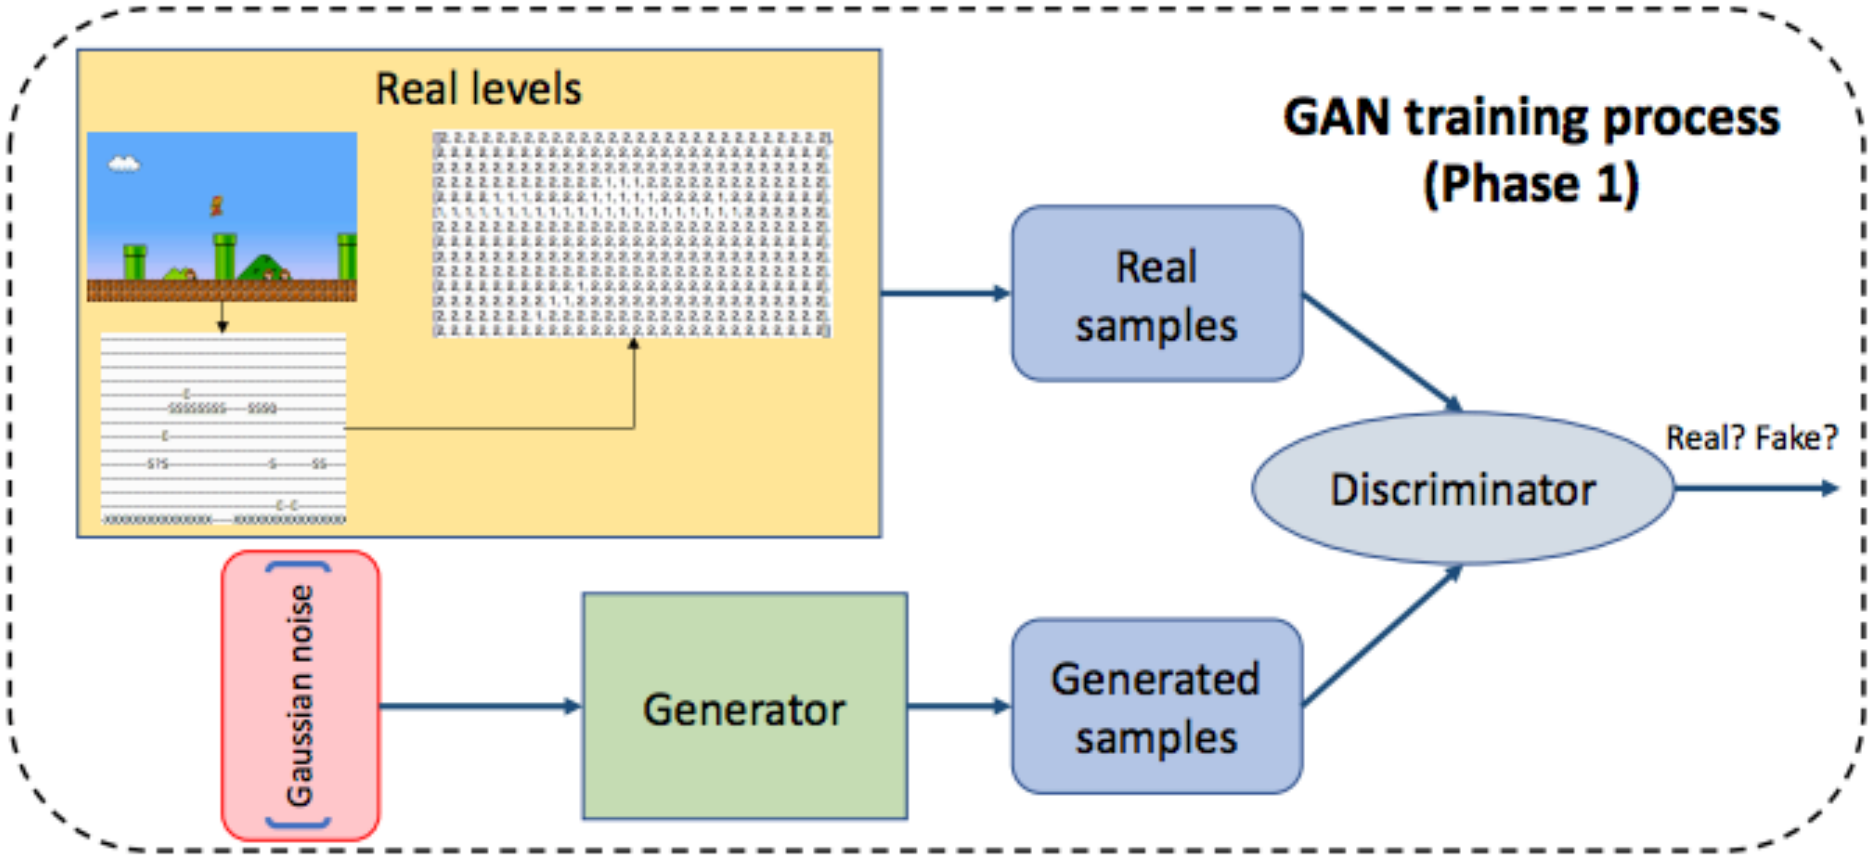

Approach – Phase II

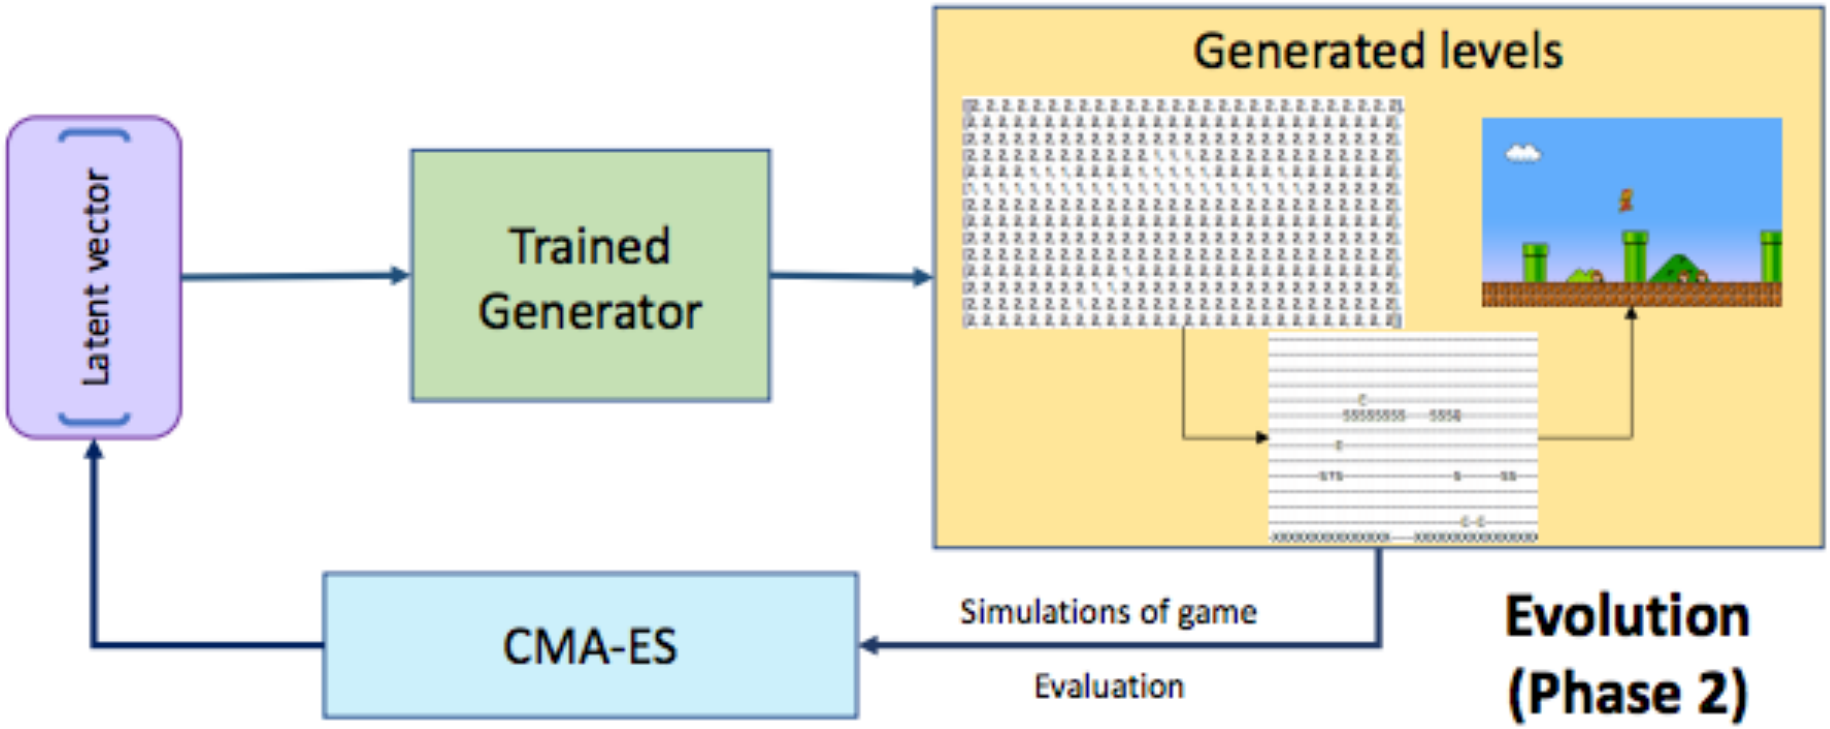

GAN Training

173 training images of size 28x14

Level Representation

Tile type	Symbol	Identity	Visualization
Solid/Ground	X	0	
Breakable	S	1	
Empty (passable)	-	2	
Full question block	?	3	
Empty question block	Q	4	
Enemy	E	5	
Top-left pipe	<	6	
Top-right pipe	>	7	
Left pipe	[	8	
Right pipe	]	9	

GAN changes:

- One-hot encoding
- ReLU activation function for output layer
- Argmax to determine tile type

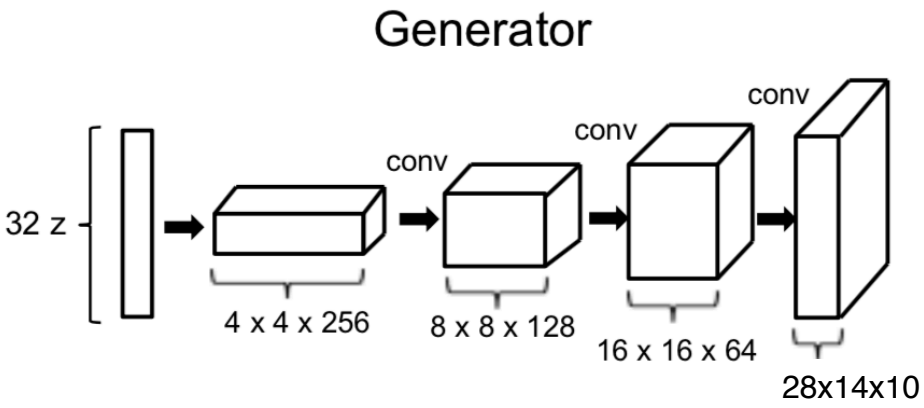

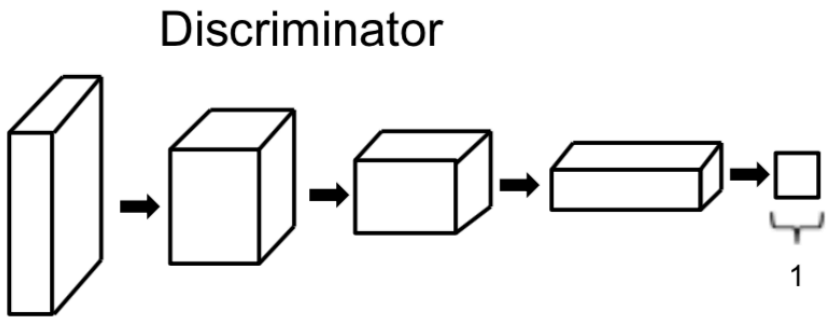

CMA-ES Experiments

- **Representation-based testing:**
  - Optimize for certain number of ground titles

$$F_{ground} = \sqrt{(g-t)^2}$$

  - Increasing difficulty (less ground, more enemies)

- **Agent-based testing:**
  A* Mario agent by Baumgarten

  Fitness = %playable + #jumps

Random Sampling

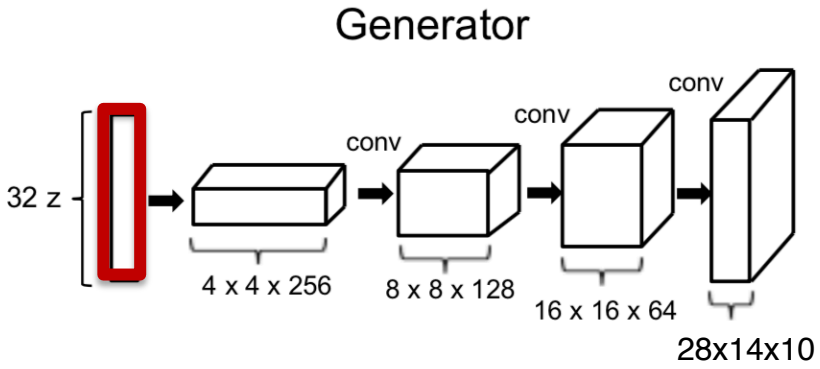

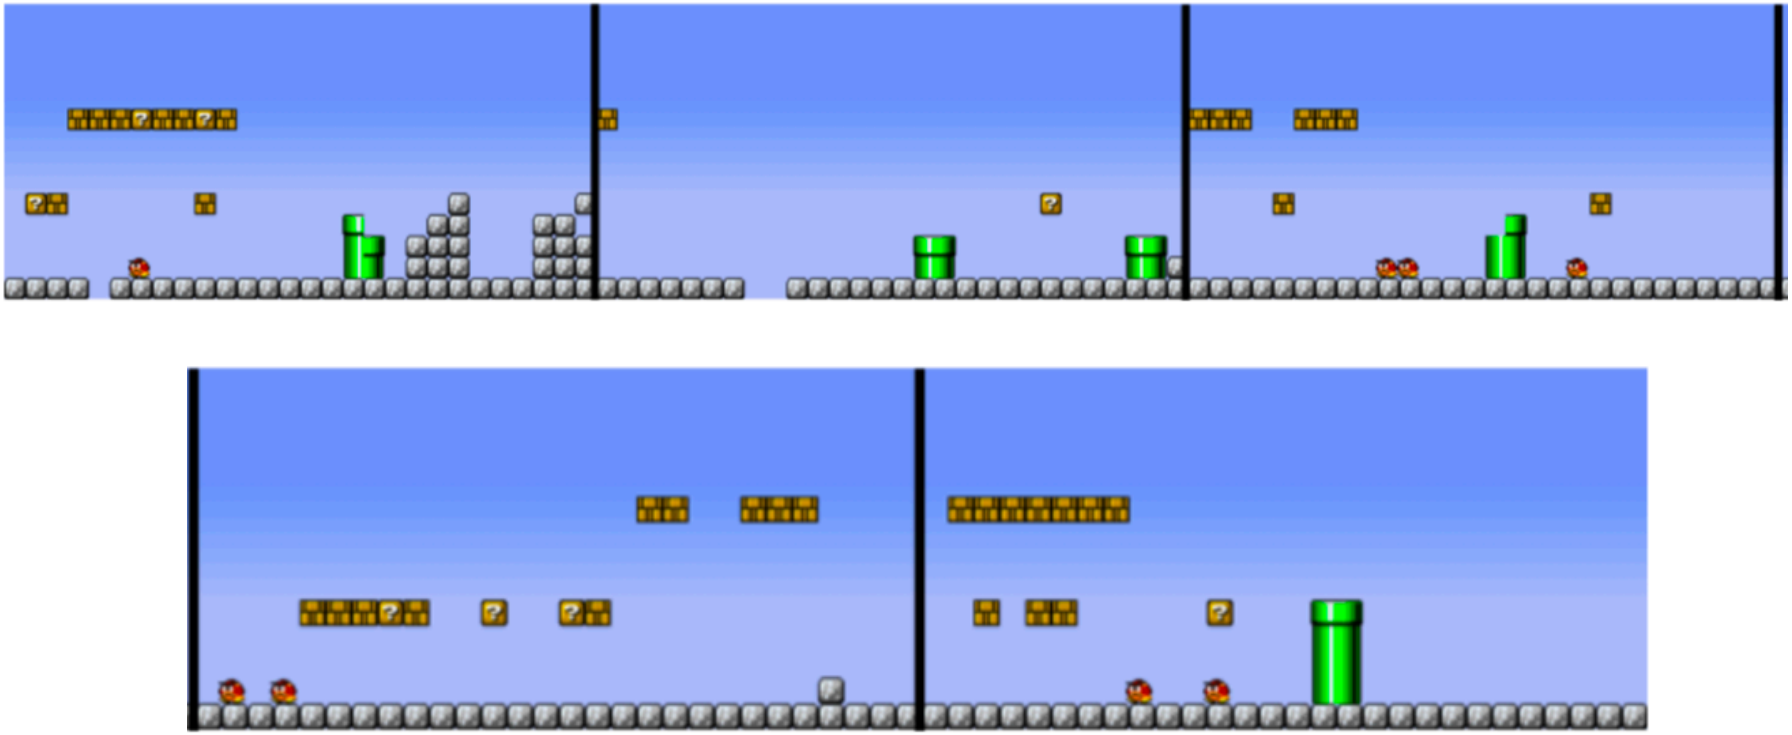

- Trained GAN can express different level variations (can be different to levels used for training)
- Captures domain regularities

Mutations

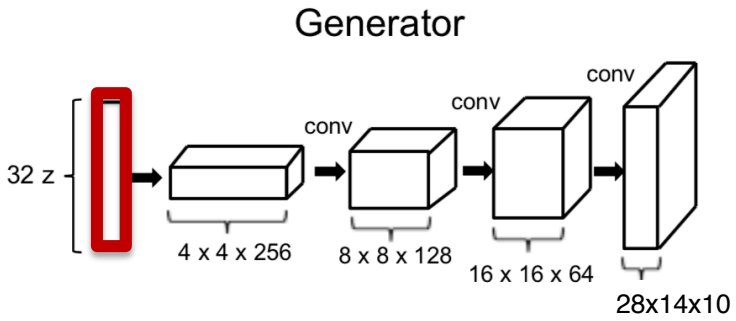

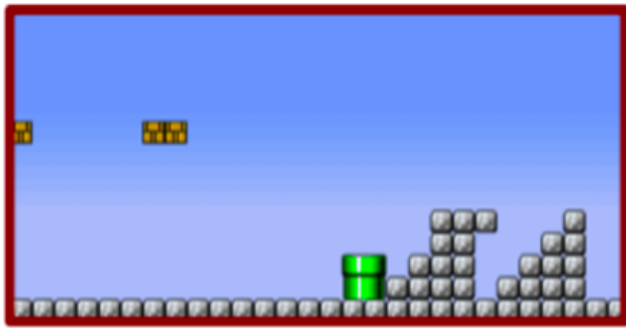

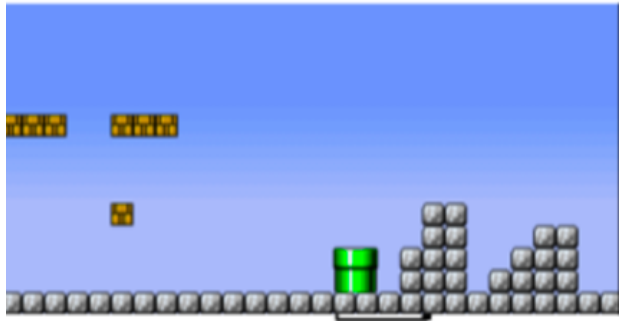

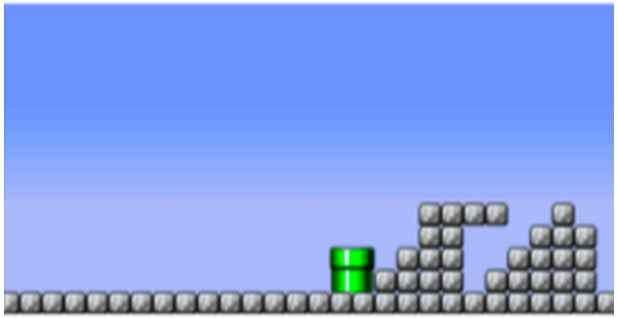

Parent

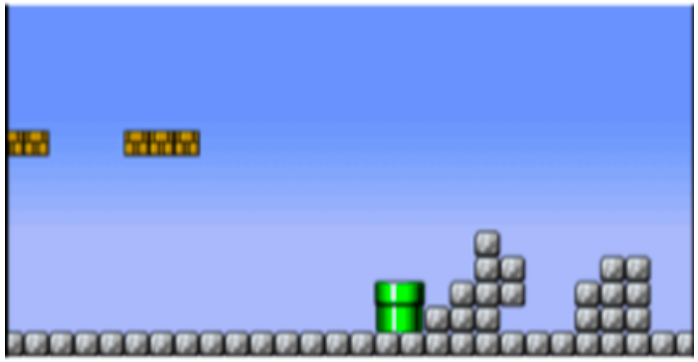

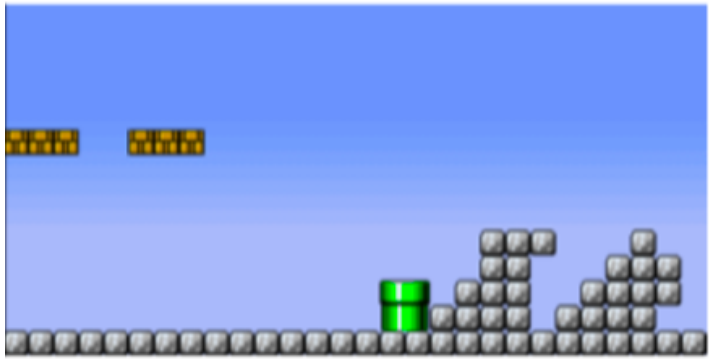

➔ Trained GAN representation displays locality

Training


Run 1

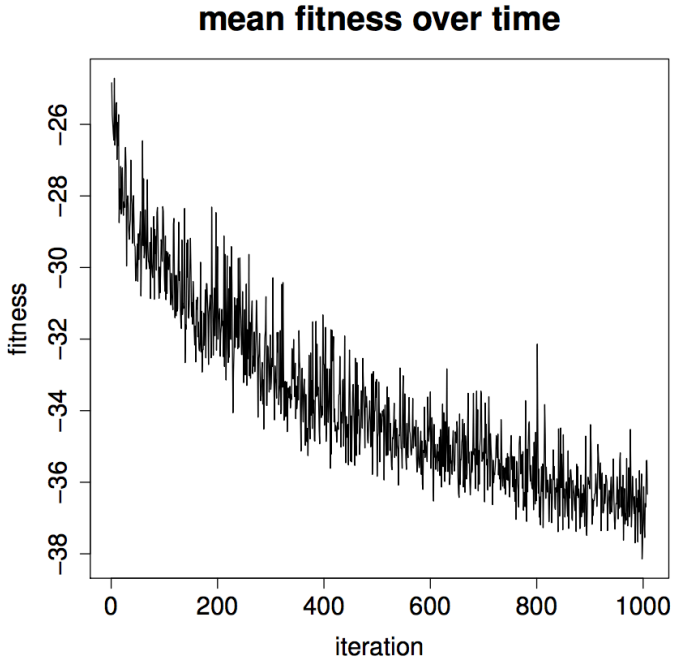
mean fitness over time
fitness
−26
−28
−30
−32
−34
−36
−38
0
200
400
600
800
1000
iteration

Results

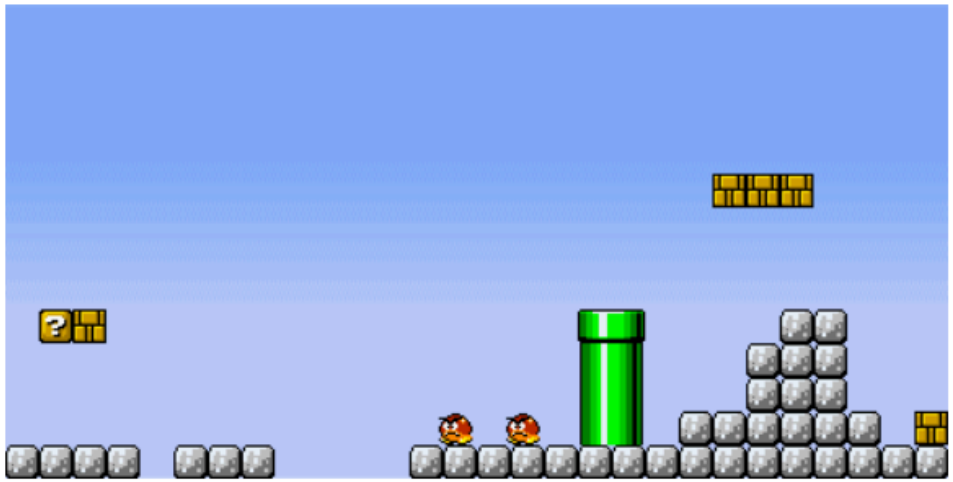

(a) Playable level maximizing jumps

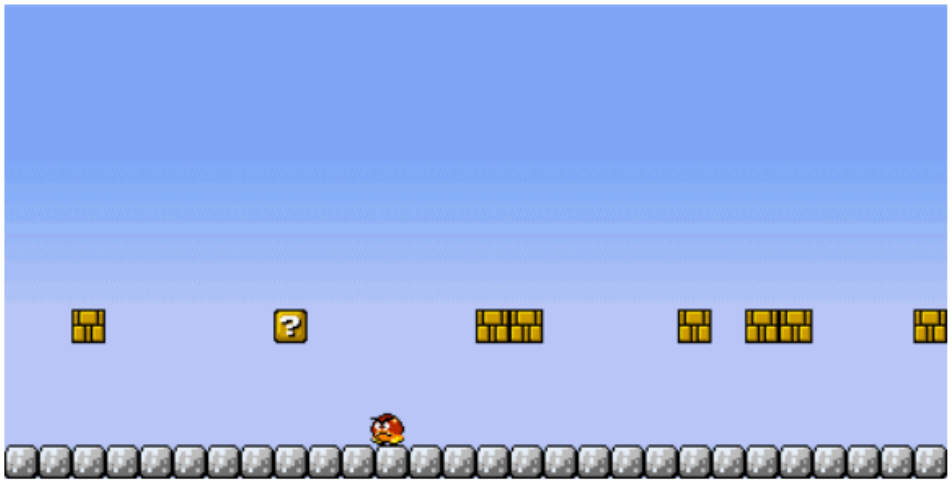

(b) Playable level minimizing jumps

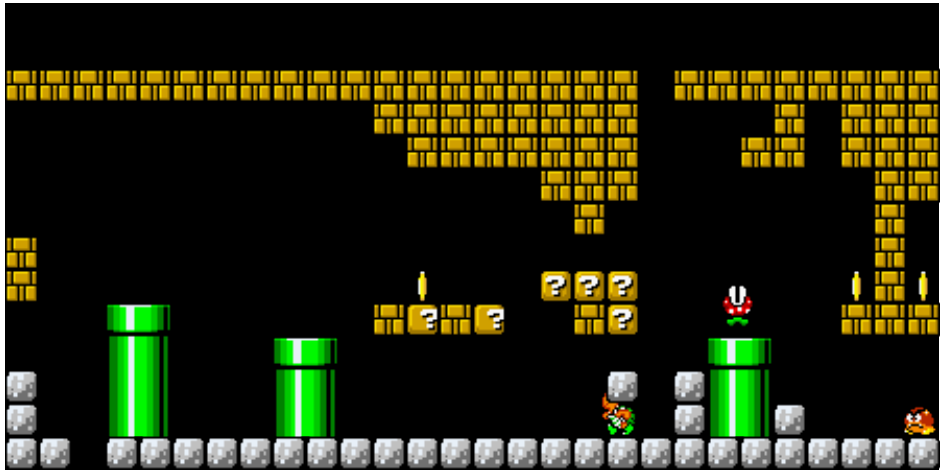

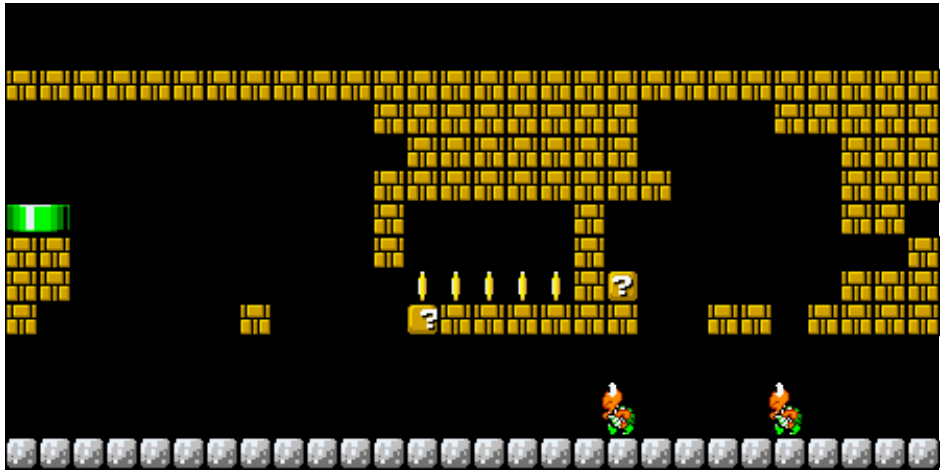

Increasing Difficulty

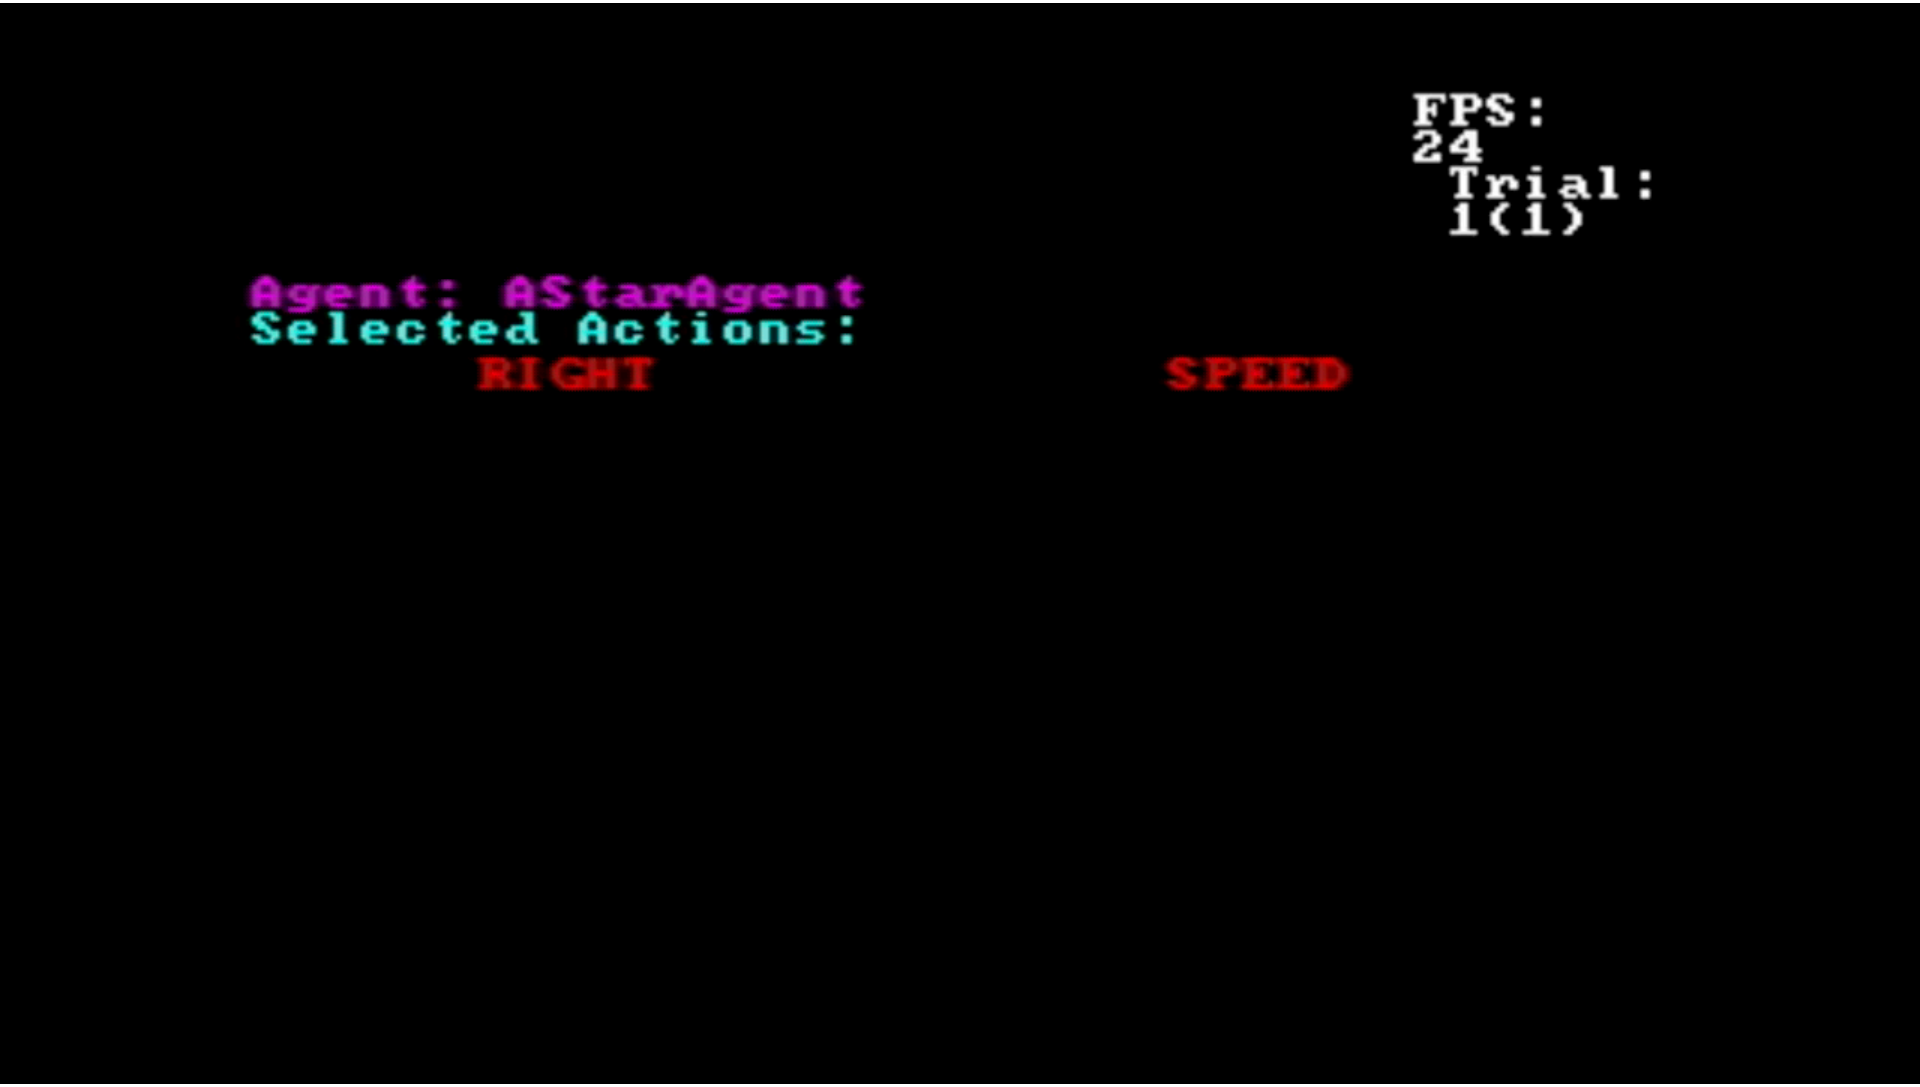

Capturing Local and Global Patterns in Procedural Content Generation via Machine Learning (CoG 2020)

Volz, Justesen, Snodgrass, Asadi, Purmonen, Holmgard, Togelius, Risi

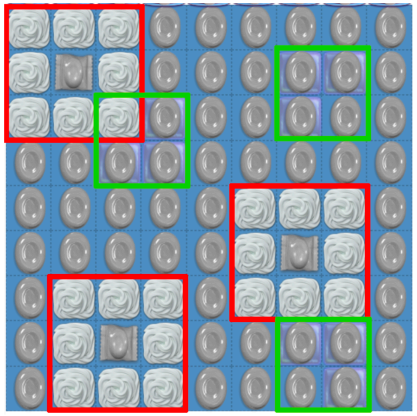

Local patterns

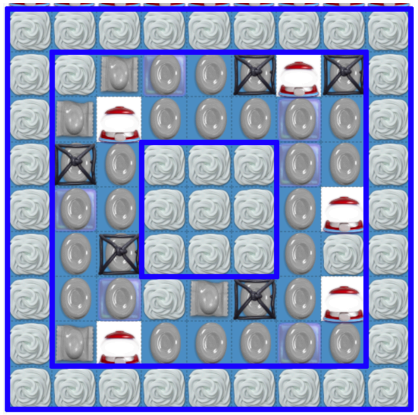

Global patterns

Towards pattern-aware PCGML:

- Enrich the data
- Augmenting the algorithm
- Filtering solutions

modl.ai

King

Results

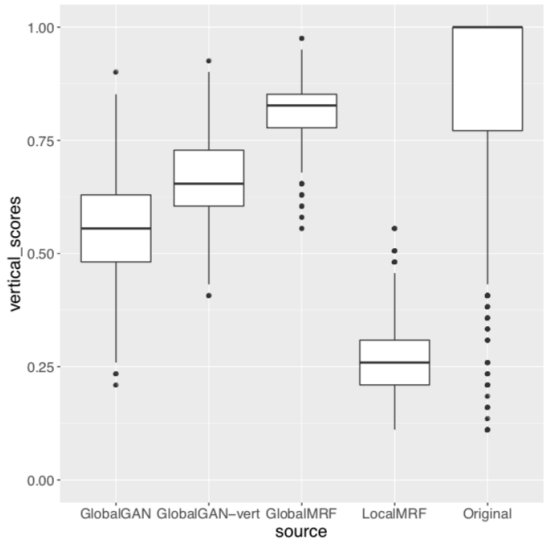

(a) Vertical symmetry.

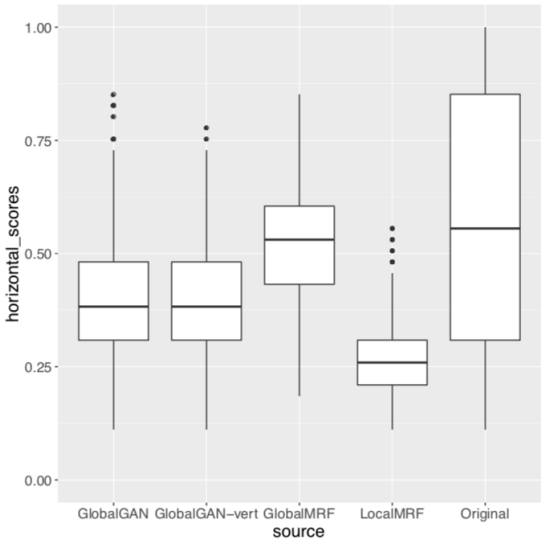

(b) Horizontal Symmetry

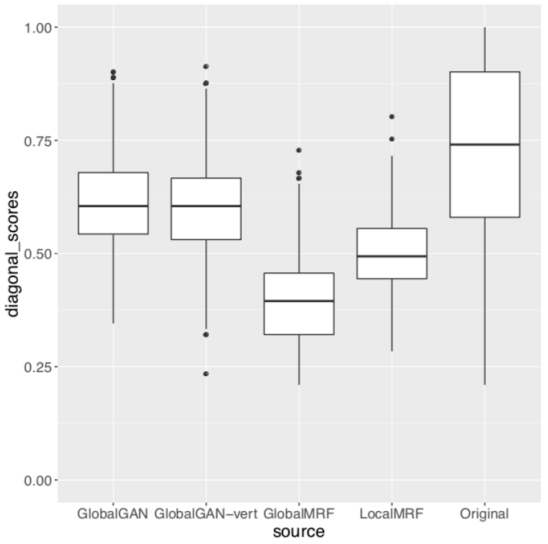

(c) Diagonal symmetry.

- Trained on 504 Candy Crush levels
- Ability to generate symmetric levels could be improved
- Other promising approaches to try:
  - CoordConv (Liu et al, 2018)
  - Self-Attention GANs (Zhang et al, 2019)


Score:
0
0 | 1
25
36
Select boosters:
Play!

What if I don’t have many levels to learn from?

Bootstrapping Conditional GANs for Video Game Level Generation

Torrado, Khalifa, Green, Justesen, Risi, Togelius, (CoG 2020)

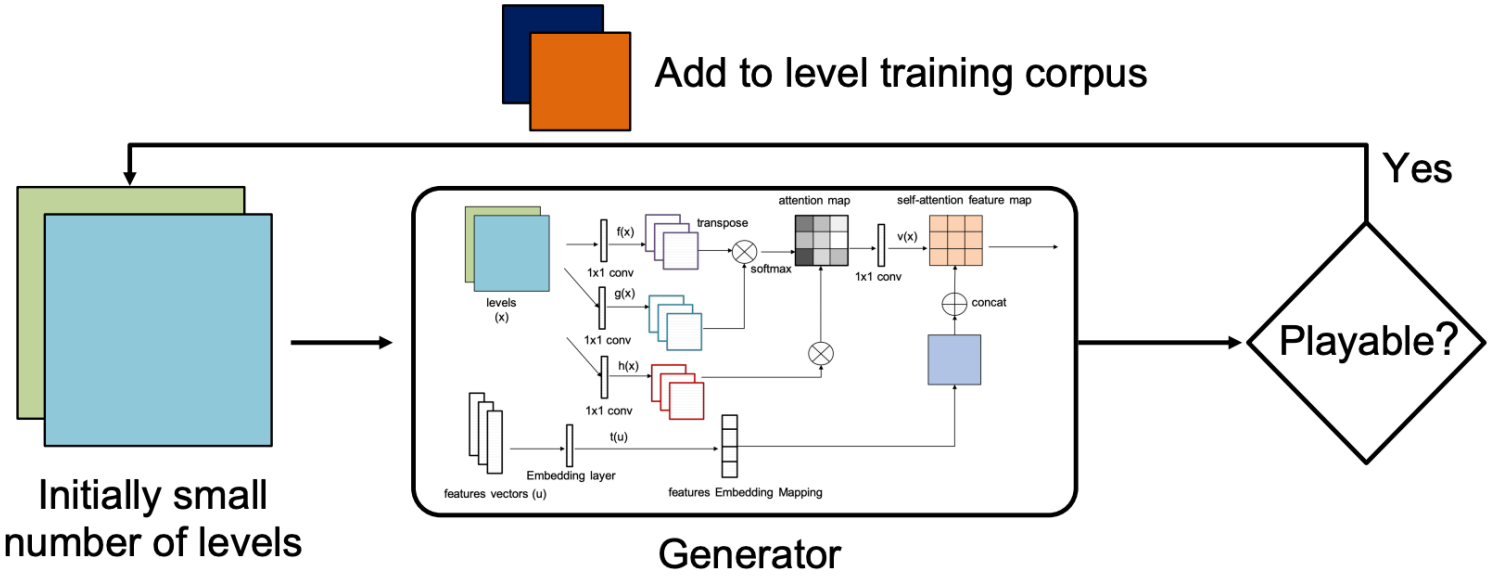

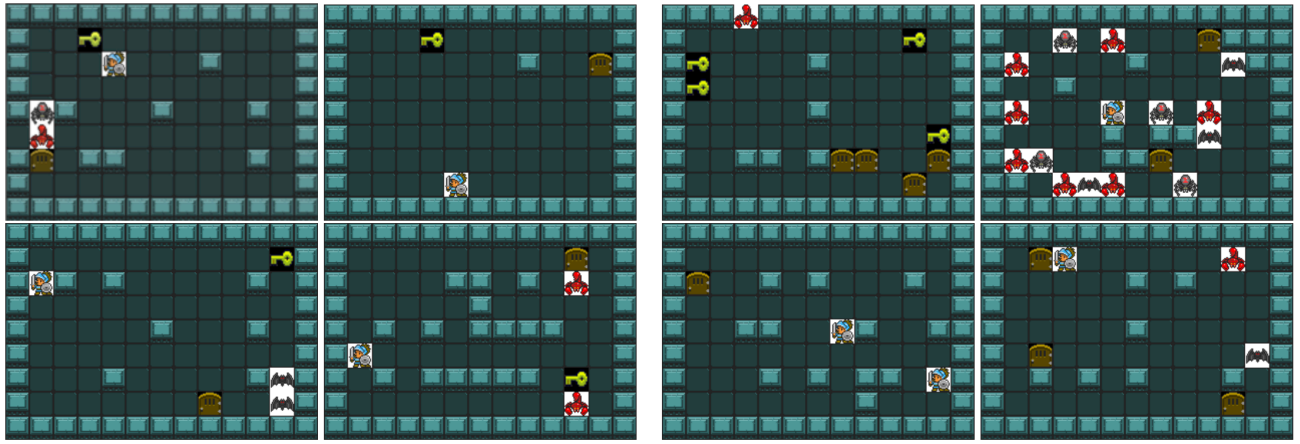

(a) Playable Levels

(b) Unplayable Levels

Interactive Latent Variable Evolution

CG-GAN: An Interactive Evolutionary GAN-based Approach for Facial Composite Generation (Zaltron, Zurlo, Risi). AAAI 2020

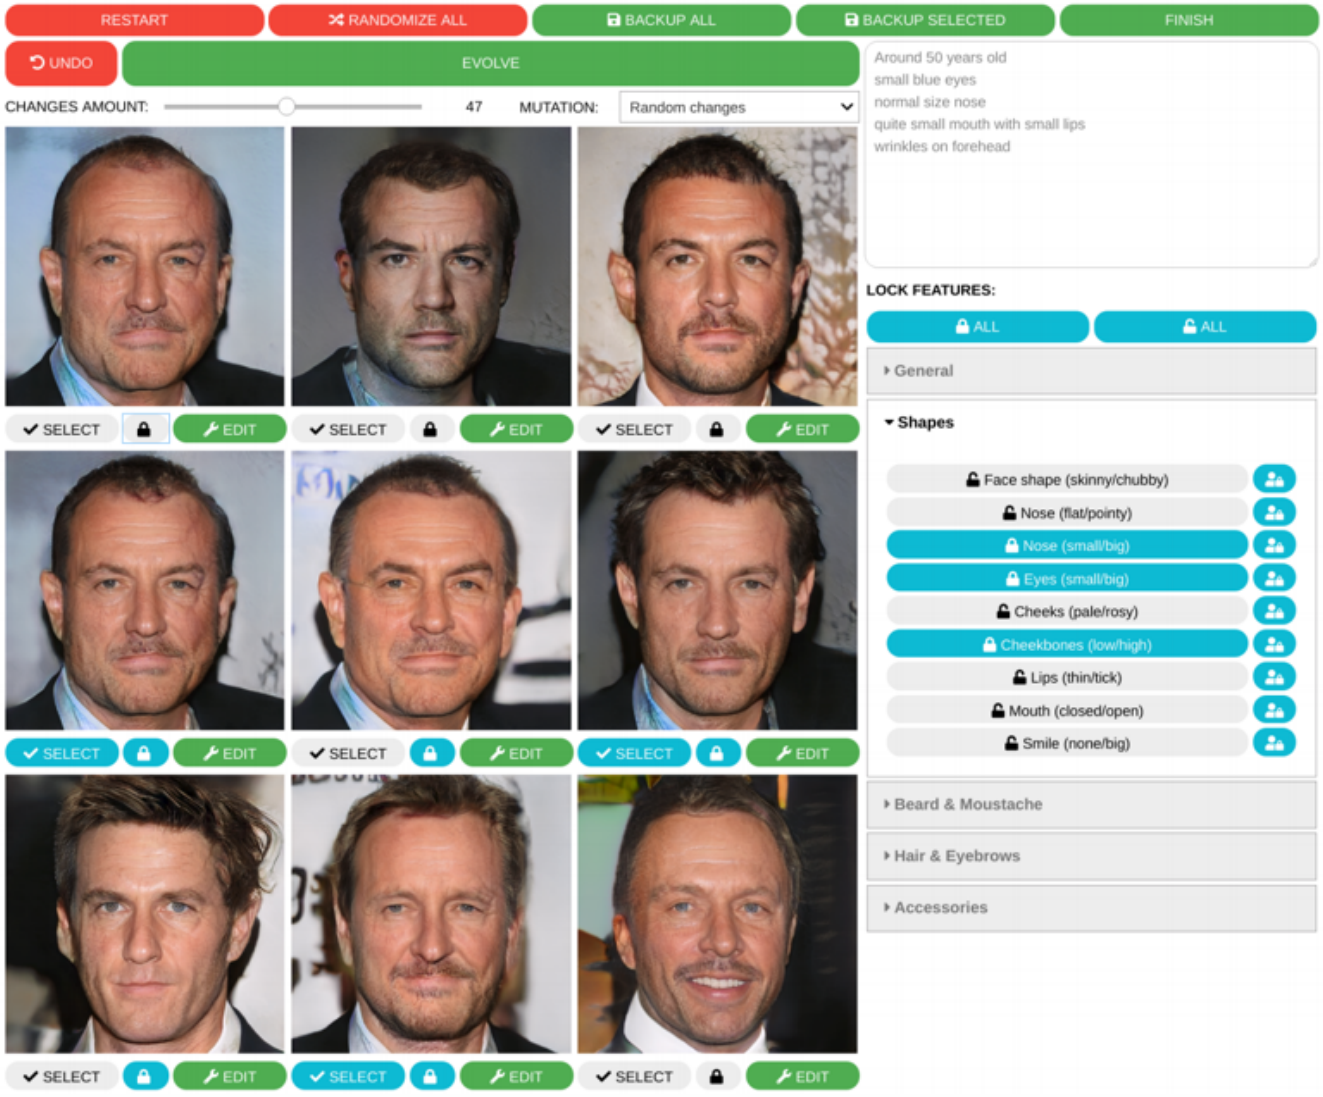

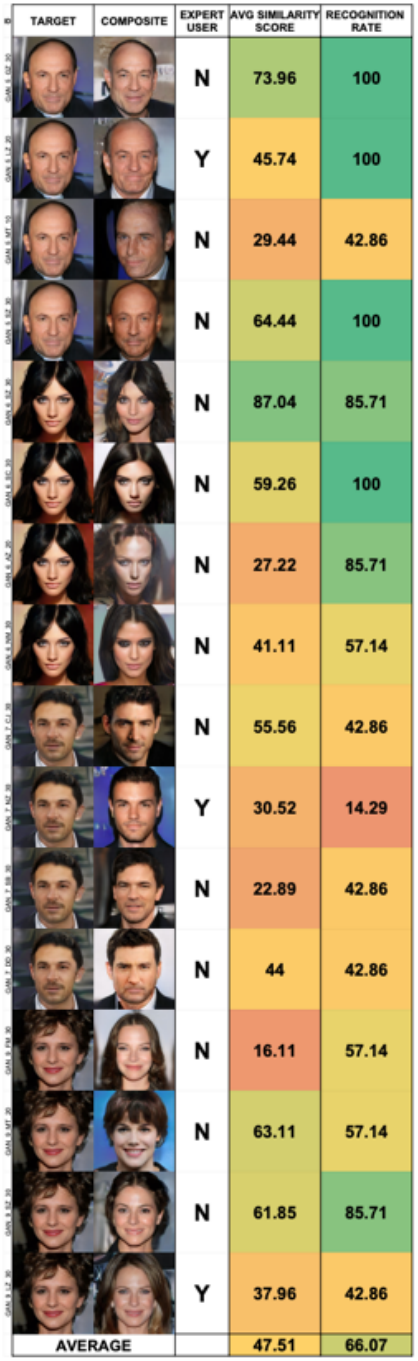

TARGET	COMPOSITE	EXPERT USER	AVG SIMILARITY SCORE	RECOGNITION RATE
		N	73.96	100
		Y	45.74	100
		N	29.44	42.86
		N	64.44	100
		N	87.04	85.71
		N	59.26	100
		N	27.22	85.71
		N	41.11	57.14
		N	55.56	42.86
		Y	30.52	14.29
		N	22.89	42.86
		N	44	42.86
		N	16.11	57.14
		N	63.11	57.14
		N	61.85	85.71
		Y	37.96	42.86
AVERAGE			47.51	66.07

Interactive Evolution and Exploration Within Latent Level-Design Space of Generative Adversarial Networks (Schrum, Gutierrez, Volz, Liu, Lucas, Risi) GECCO 2020

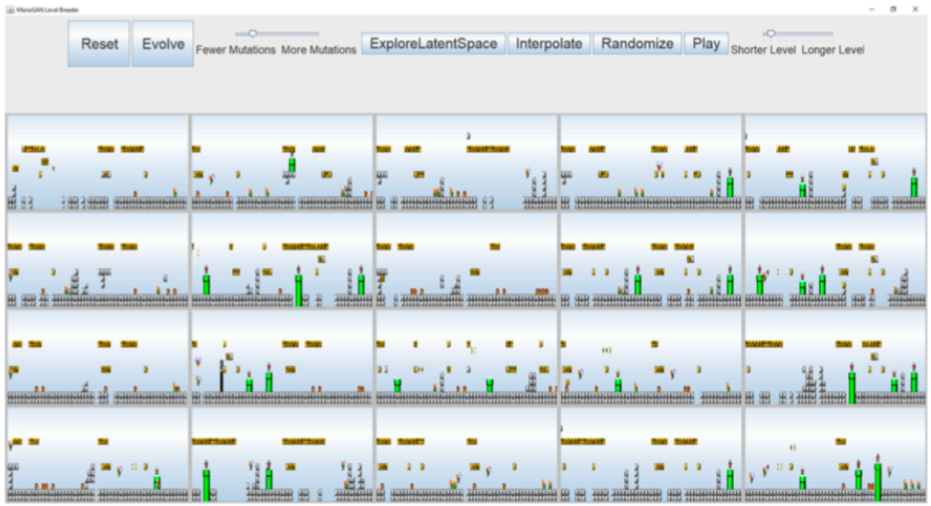

(a) Super Mario Bros.

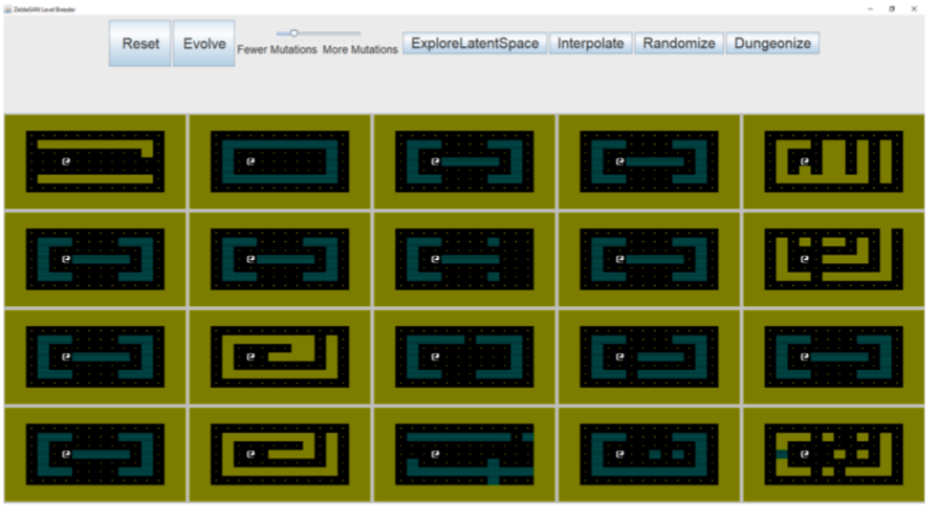

(b) The Legend of Zelda

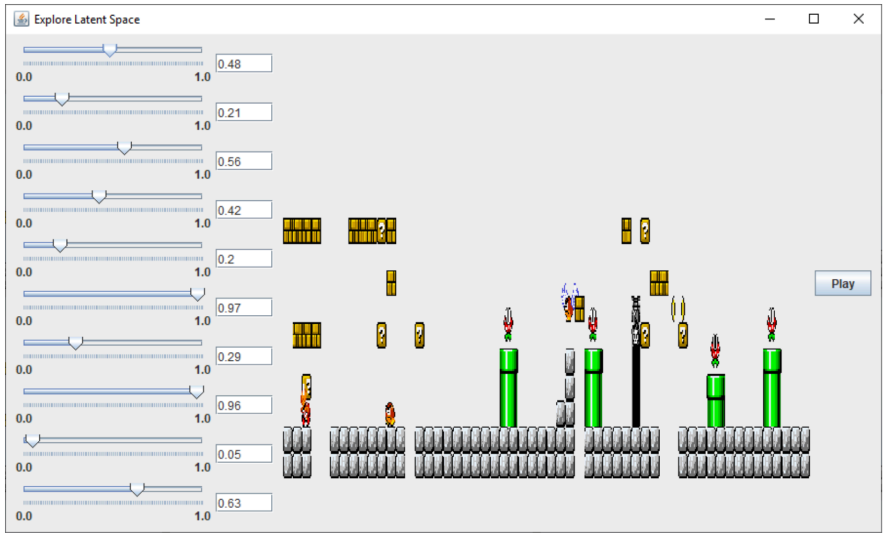

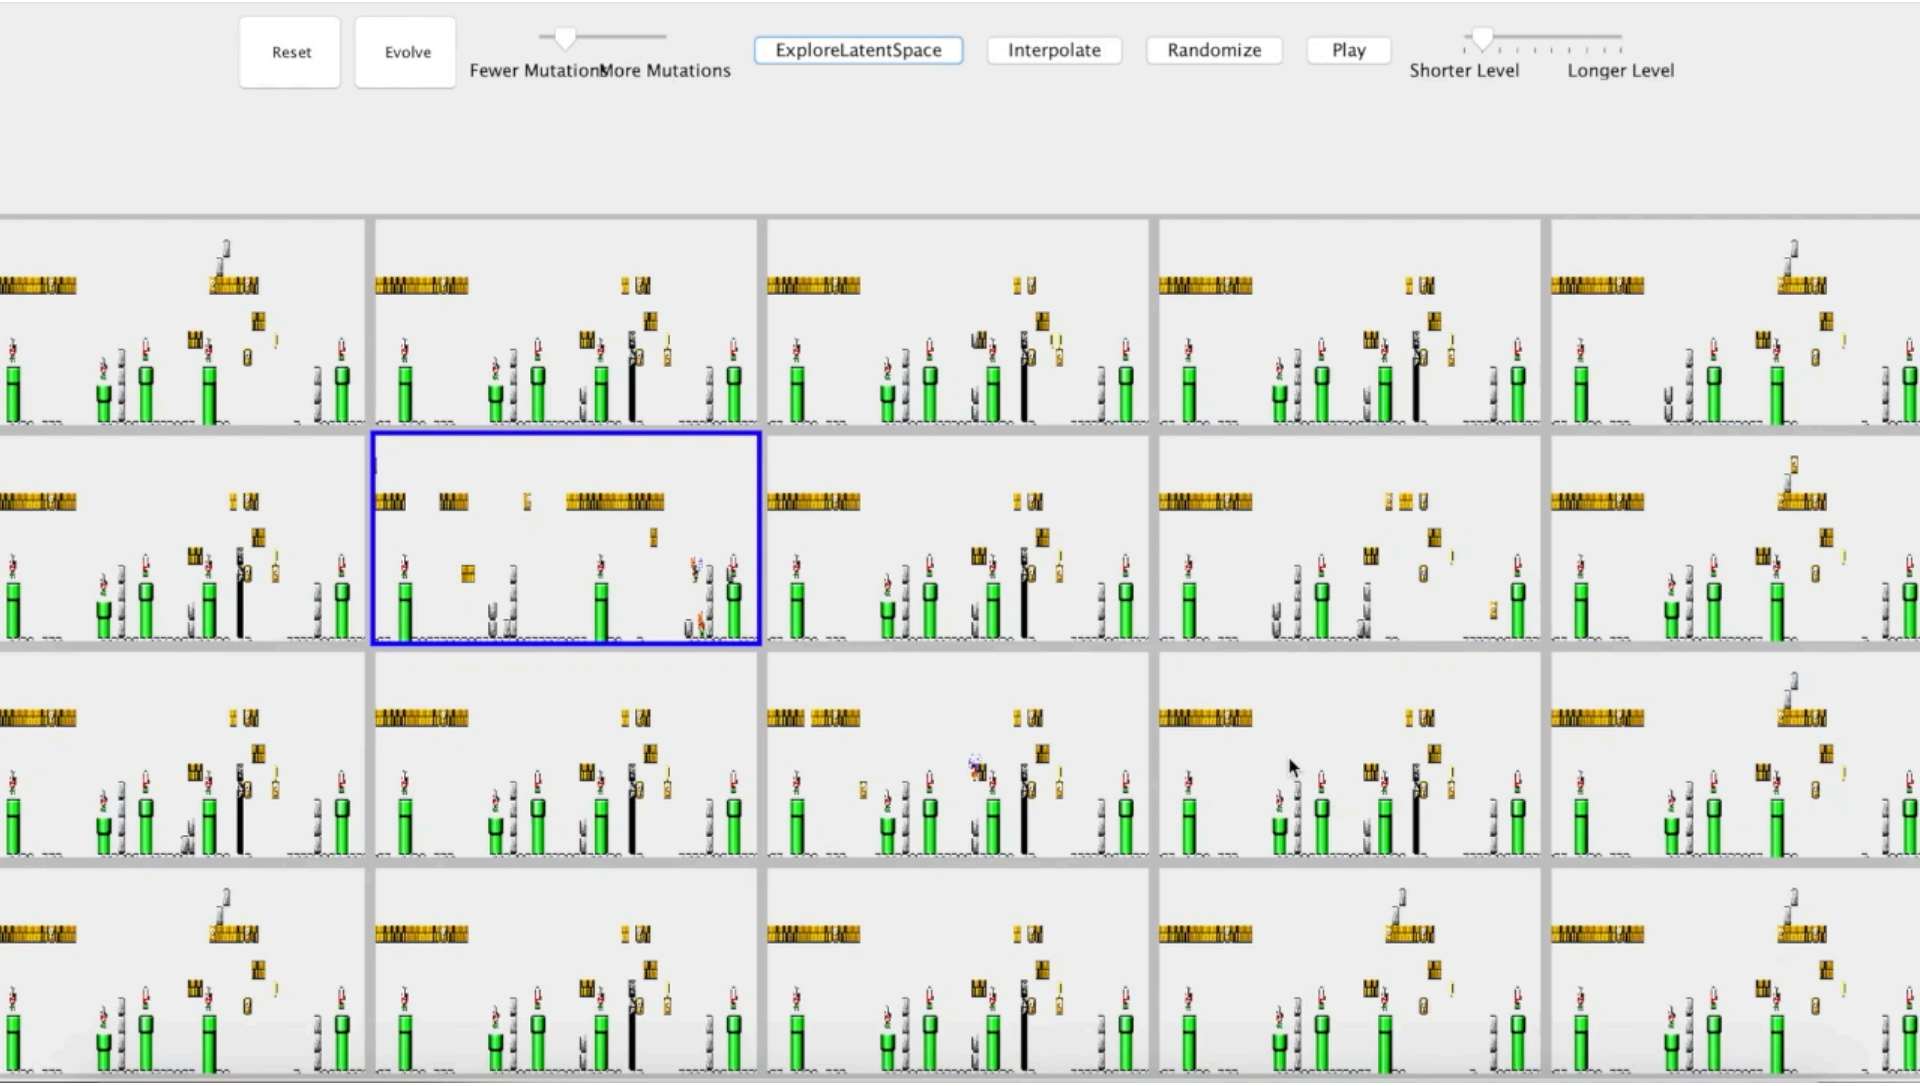
Reset
Evolve
Fewer Mutations
More Mutations
ExploreLatentSpace
Interpolate
Randomize
Play
Shorter Level
Longer Level

How can we scale this to larger patterns?

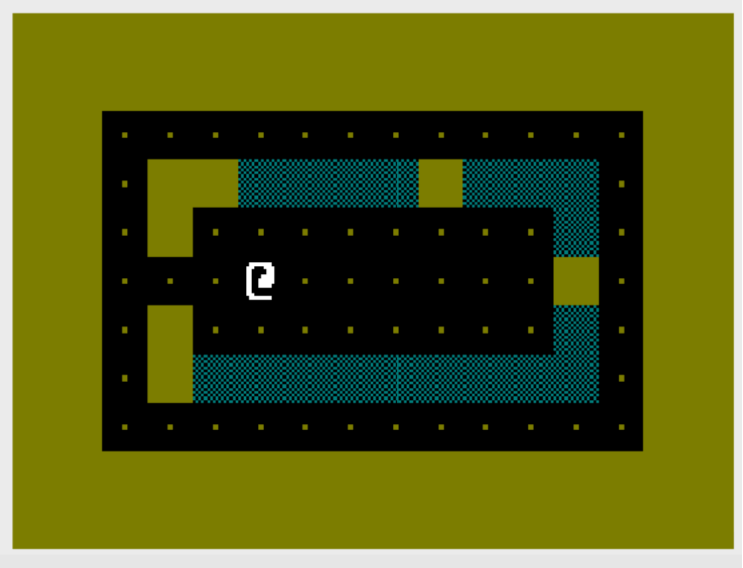

CPPN2GAAN

Schrum, Volz, Risi (GECCO 2020)

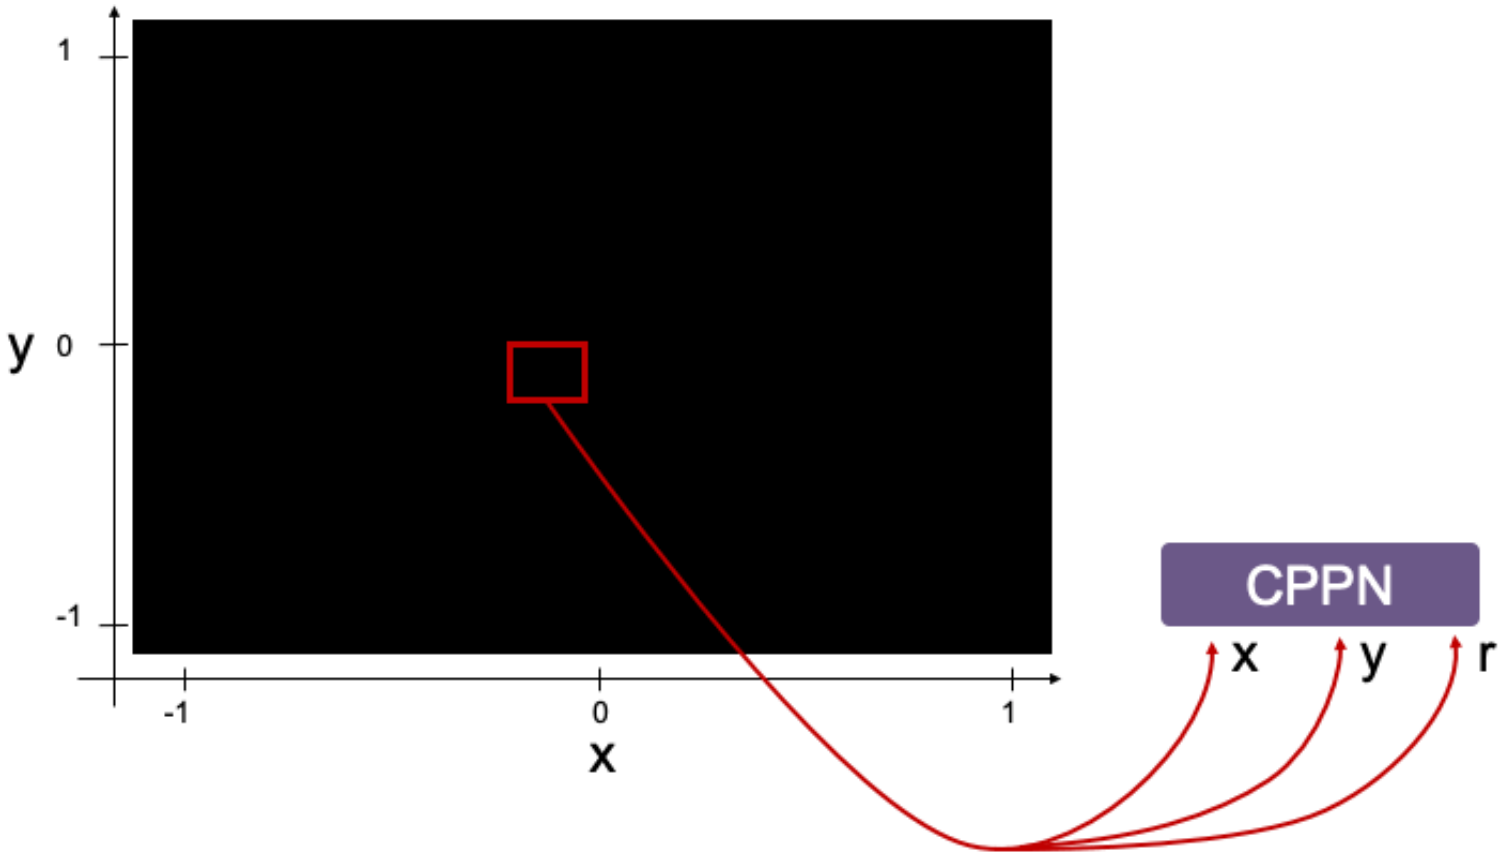

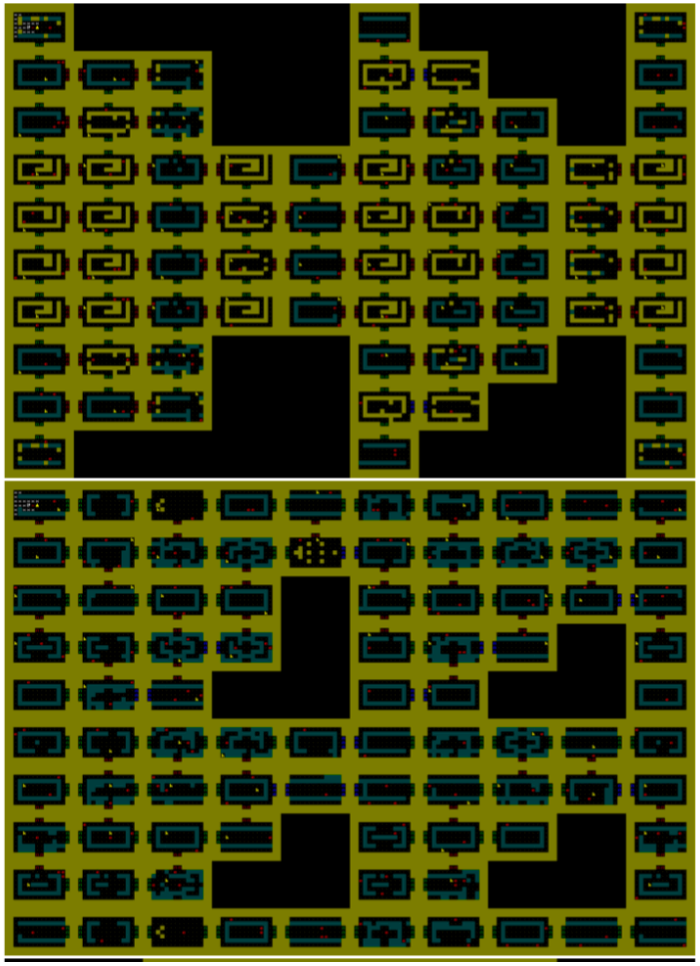

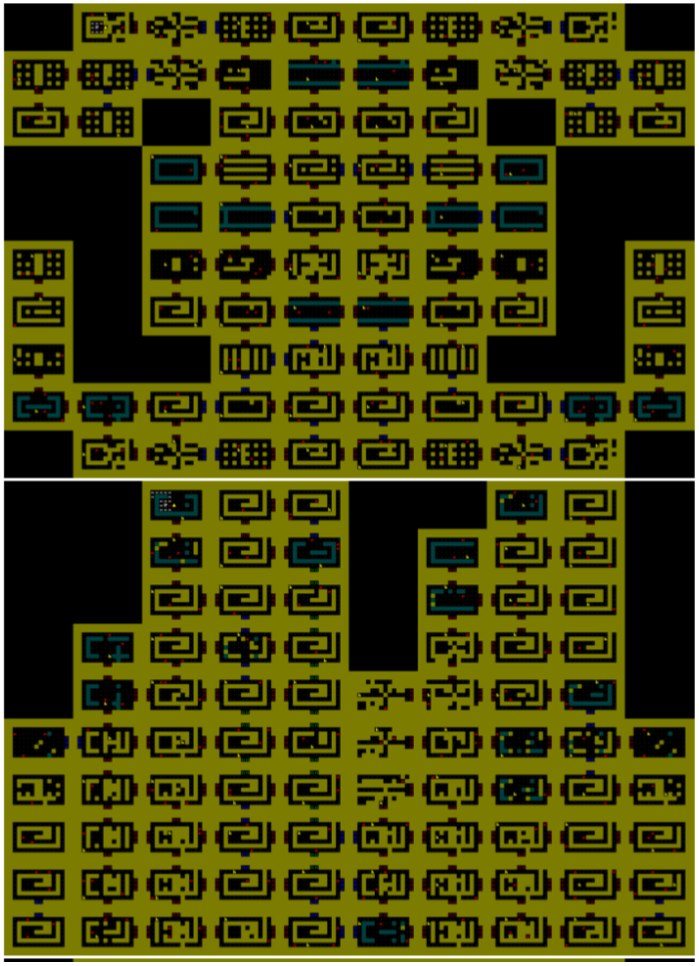

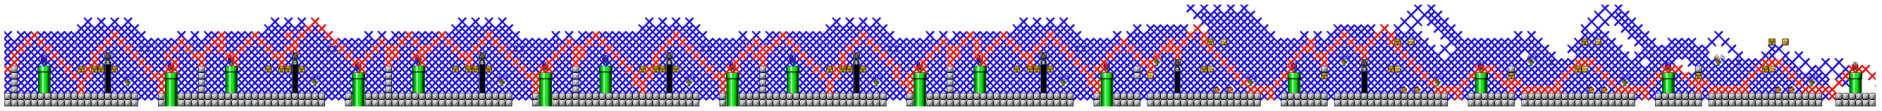

MAP-Elites Coverage

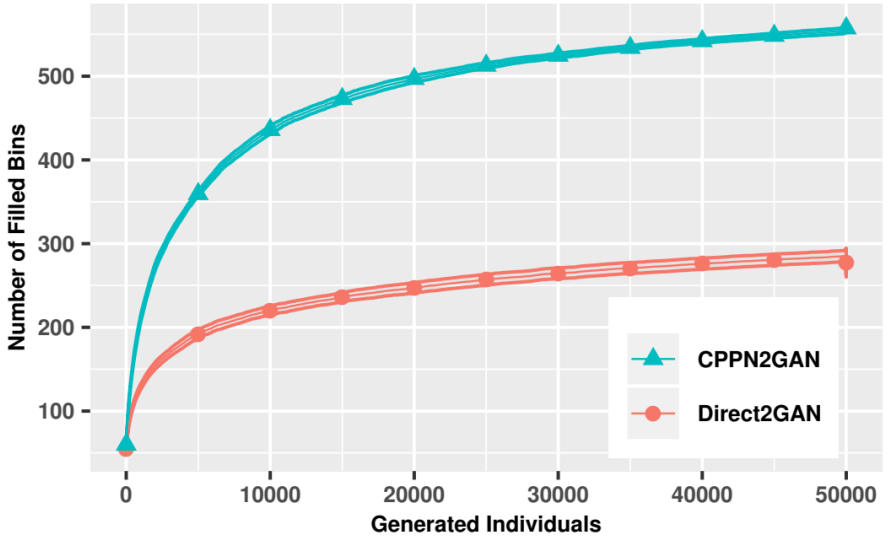

(a) Mario

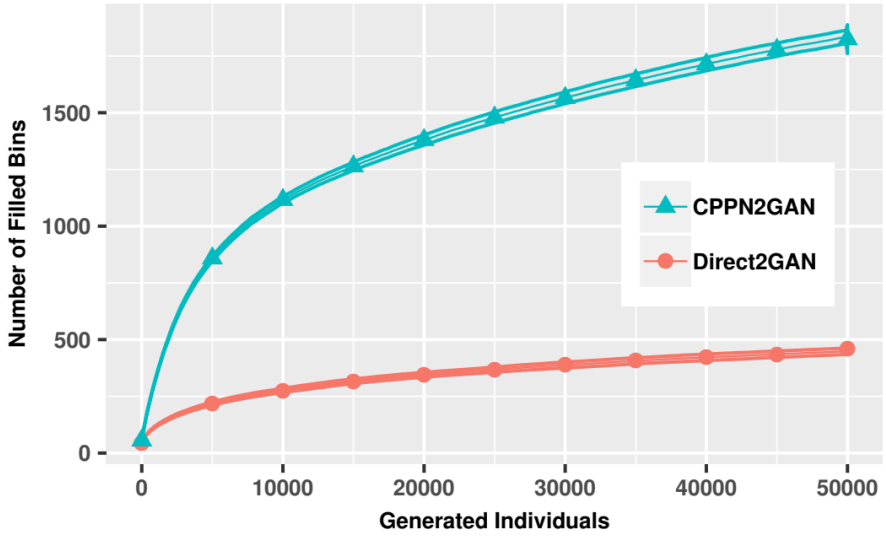

(b) Zelda

I’m not interested in procedural content generation, why should I care?

Discovering Representations for Black-box Optimization, Gaier, Asteroth, Mouret (GECCO 2020)

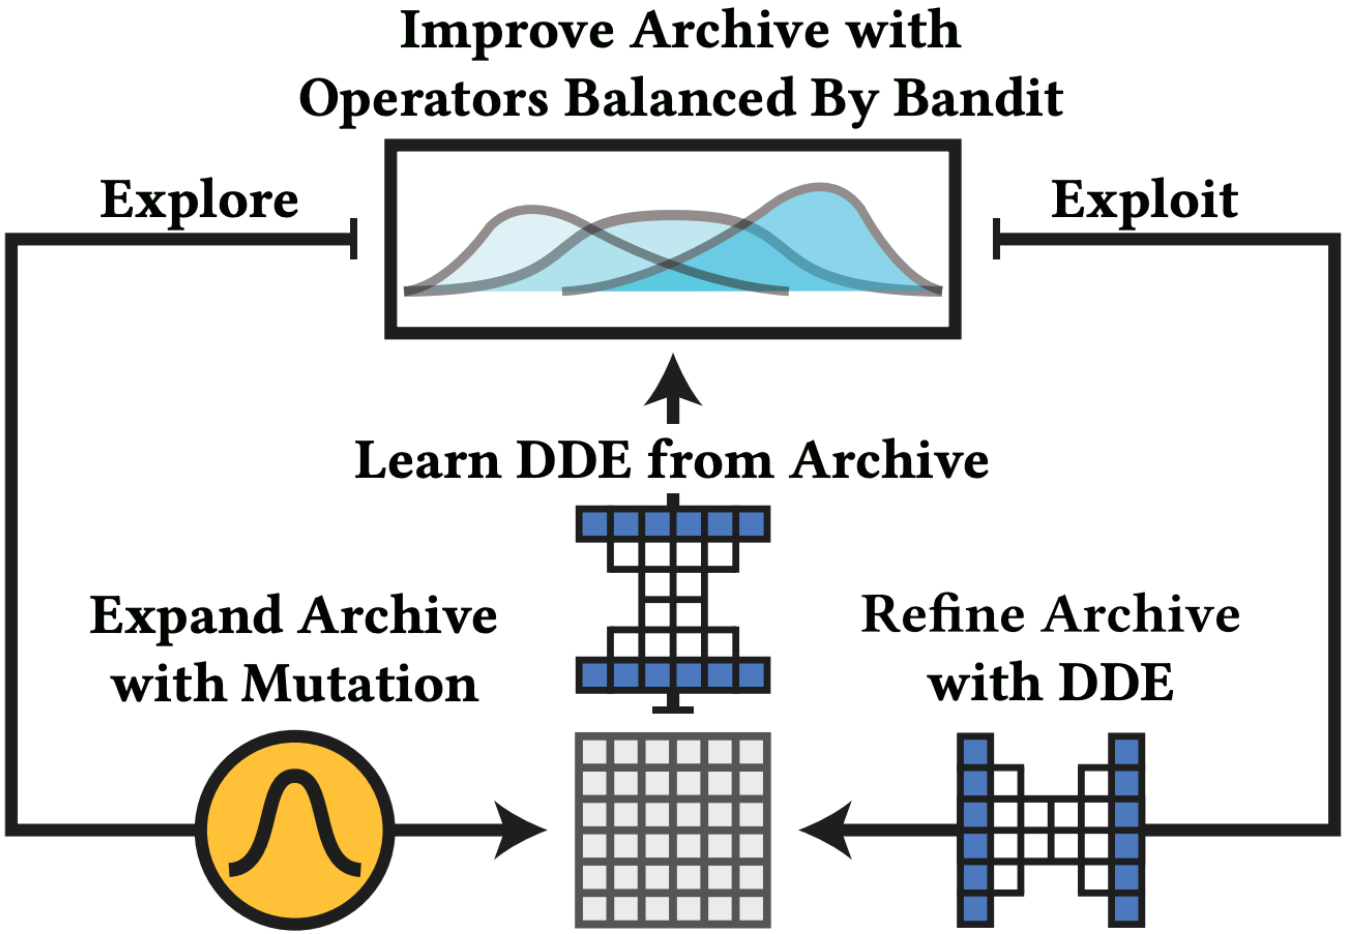

Conclusion

- Data-driven encodings a promising approach for EC
- Can also be combined with other representations such as CPPNs
- Next: Data-driven encodings for control tasks

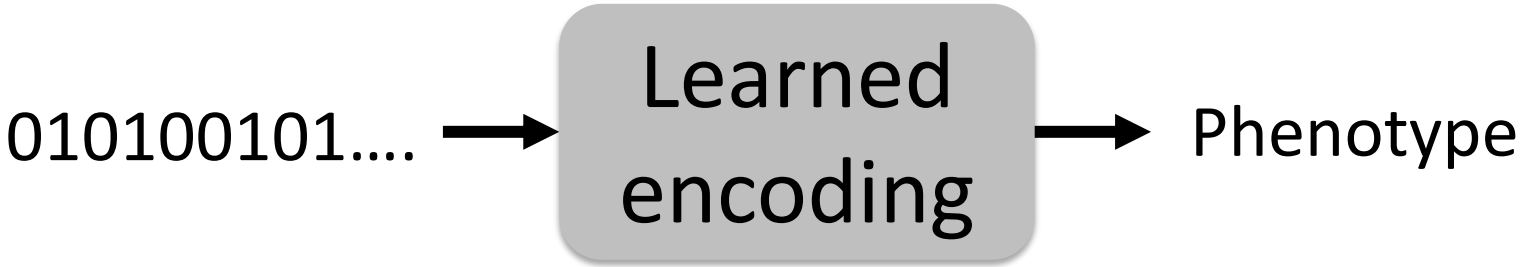

Thank you for your attention! Questions?

- www.sebastianrisi.com
- Email: sebr@itu.dk

 @risi1979