

Generative AI Models

ECE 598 LV – Lecture 1

Lav R. Varshney

18 January 2022

GAN PROGRESS ON FACE GENERATION

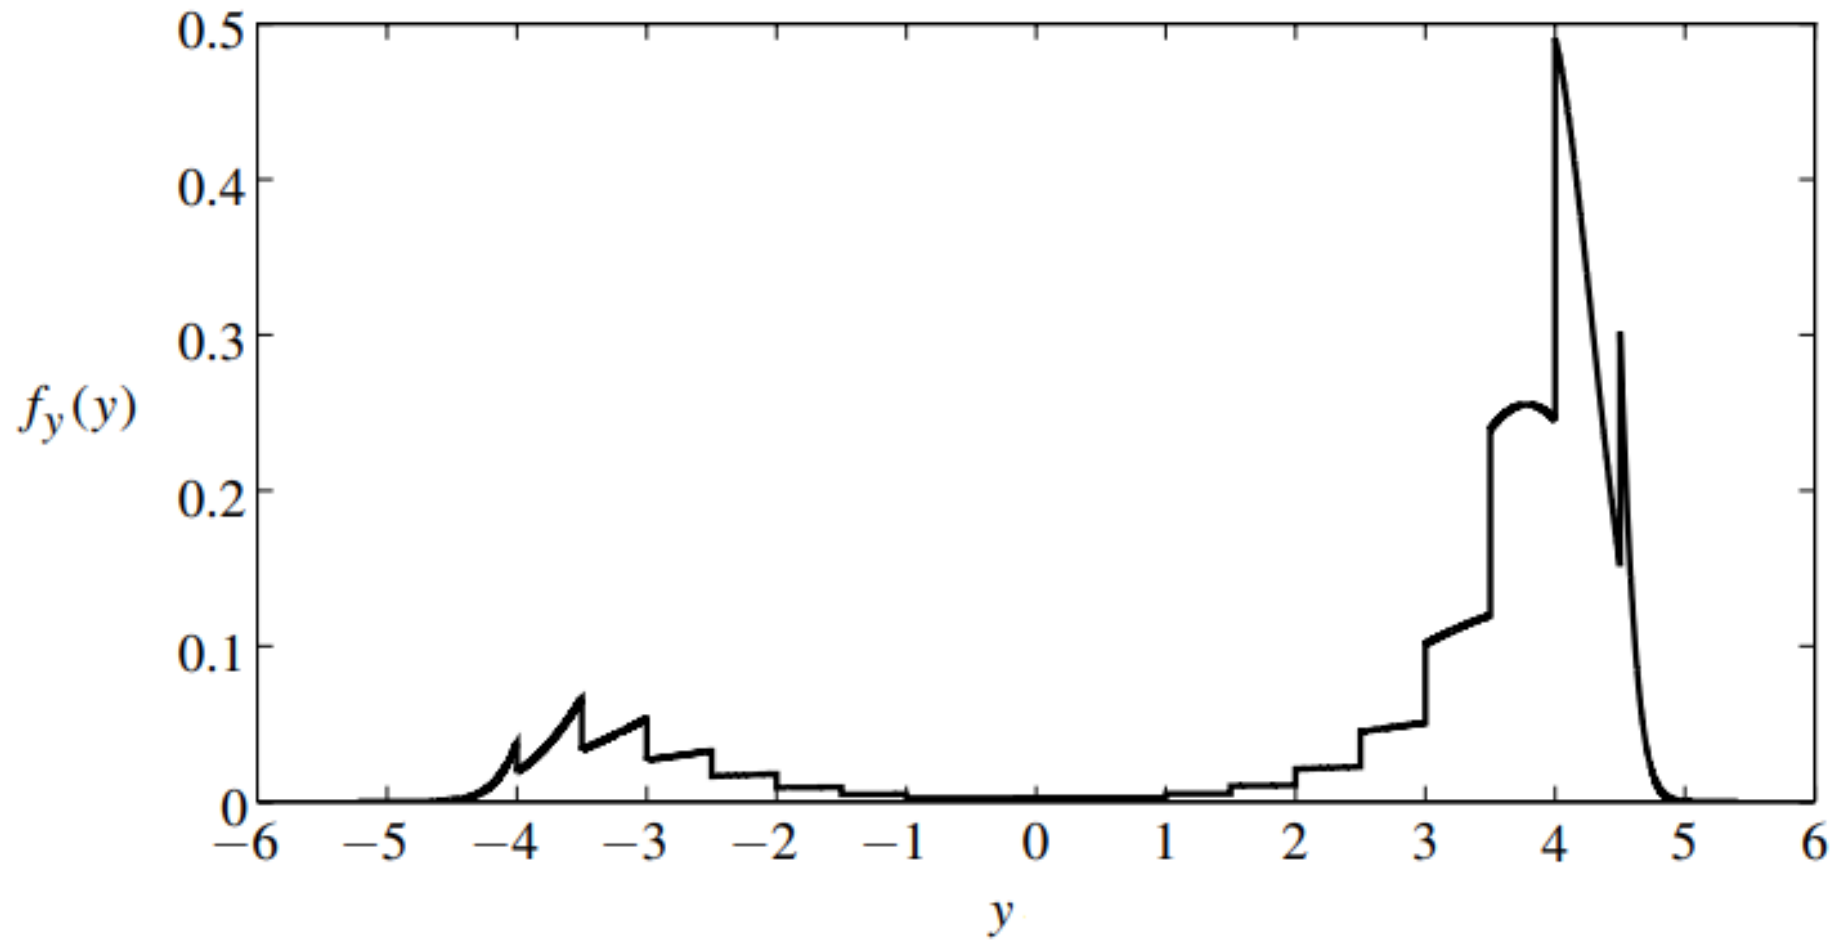
Source: Goodfellow et al., 2014; Radford et al., 2016; Liu & Tuzel, 2016; Karras et al., 2018; Karras et al., 2019; Goodfellow, 2019; Karras et al., 2020; AI Index, 2021



Figure 2.1.7

Links <https://www.cnn.com/2007/09/20/us-president-meets-british-pm>
LONDON, England (CNN) – U.S. President George W. Bush met with British Prime Minister Tony Blair on Monday to discuss the war in Iraq, according to a statement from Blair’s office.
The meeting was held at 10 Downing Street and lasted about an hour.
”The two leaders discussed Iraq and other international issues of mutual concern,” said Blair.
Bush, who is scheduled to meet Wednesday with Russian President Vladimir Putin, will also visit Germany for talks later this week.
In his statement, Blair said, ”We agreed that we should continue our efforts together to bring peace and stability to Iraq. We both reaffirmed our commitment to working closely together, as well as to continuing to work constructively toward achieving lasting security and prosperity throughout the Middle East region.”
Bush’s trip comes after he visited Britain last week where he spoke out against terrorism while visiting Buckingham Palace.
He has been criticized by some lawmakers over what they say are insufficient military resources being devoted to fighting terrorism.





Bernard Widrow and István Kollár, *Quantization Noise: Roundoff Error in Digital Computation, Signal Processing, Control, and Communications*, Cambridge University Press, 2008.

Simulation

- A lot of computing software allows you to produce random (or at least pseudorandom) samples according to a uniform $[0,1]$ random variable.
- What if we want to generate random samples from some other random variable?
- Probability integral transformation
 - Applying F^{-1} to a uniform random variable should produce a random variable with cdf F

Simulation

- Let F be a function satisfying the three properties required of a cdf, and let U be uniformly distributed over the interval $[0,1]$
- The problem is to find a function g so that F is the cdf of $g(U)$
- An appropriate function g is given by the inverse function of F ,
$$F^{-1}(u) = \min\{c: F(c) \geq u\}$$

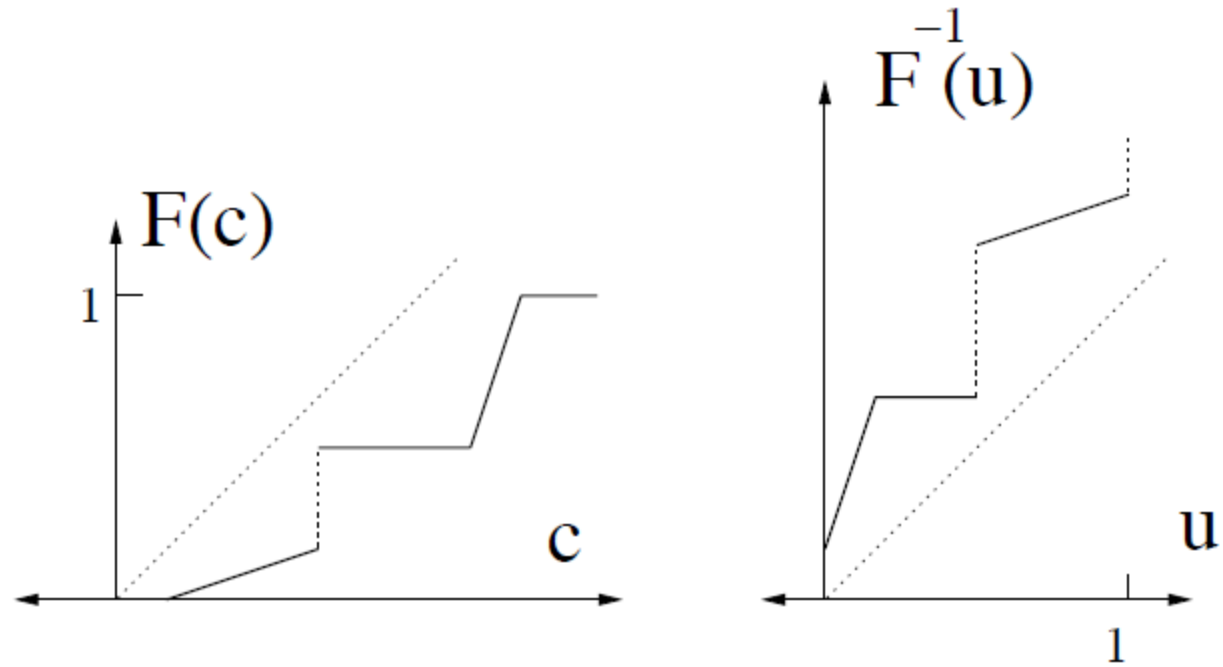
Given any random continuous variable X , define $Y = F_X(X)$.

Then:

$$\begin{aligned} F_Y(y) &= \mathbb{P}(Y \leq y) \\ &= \mathbb{P}(F_X(X) \leq y) \\ &= \mathbb{P}(X \leq F_X^{-1}(y)) \\ &= F_X(F_X^{-1}(y)) \\ &= y \end{aligned}$$

F_Y is just the CDF of a Uniform(0,1) random variable. Thus, Y has a uniform distribution on the interval $[0,1]$.

Simulation



Simulation: Example

- Simulate exponentially distributed random variable with parameter $\lambda = 1$
- $F_X(u) = 1 - e^{-u}$ for $u \geq 0$
- $F_X^{-1}(c) = -\ln(1 - c)$

```
x = rand(1,1000);
```

```
y = -log(1-x);
```

```
hist(y,100);
```

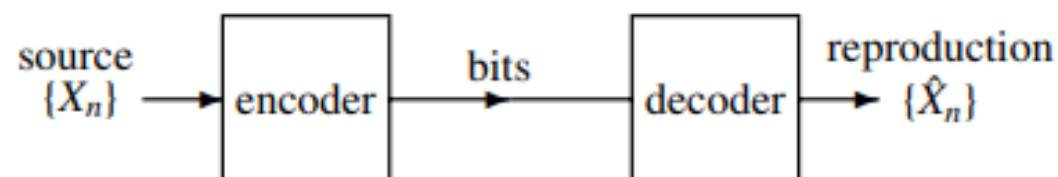
Source Coding and Simulation

XXIX Shannon Lecture, presented at the 2008 IEEE International Symposium on Information Theory, Toronto Canada

Robert M. Gray



Source coding/compression/quantization



Simulation/synthesis/fake process

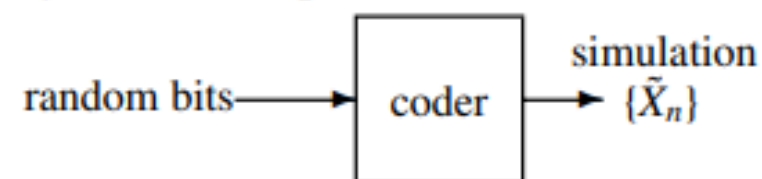
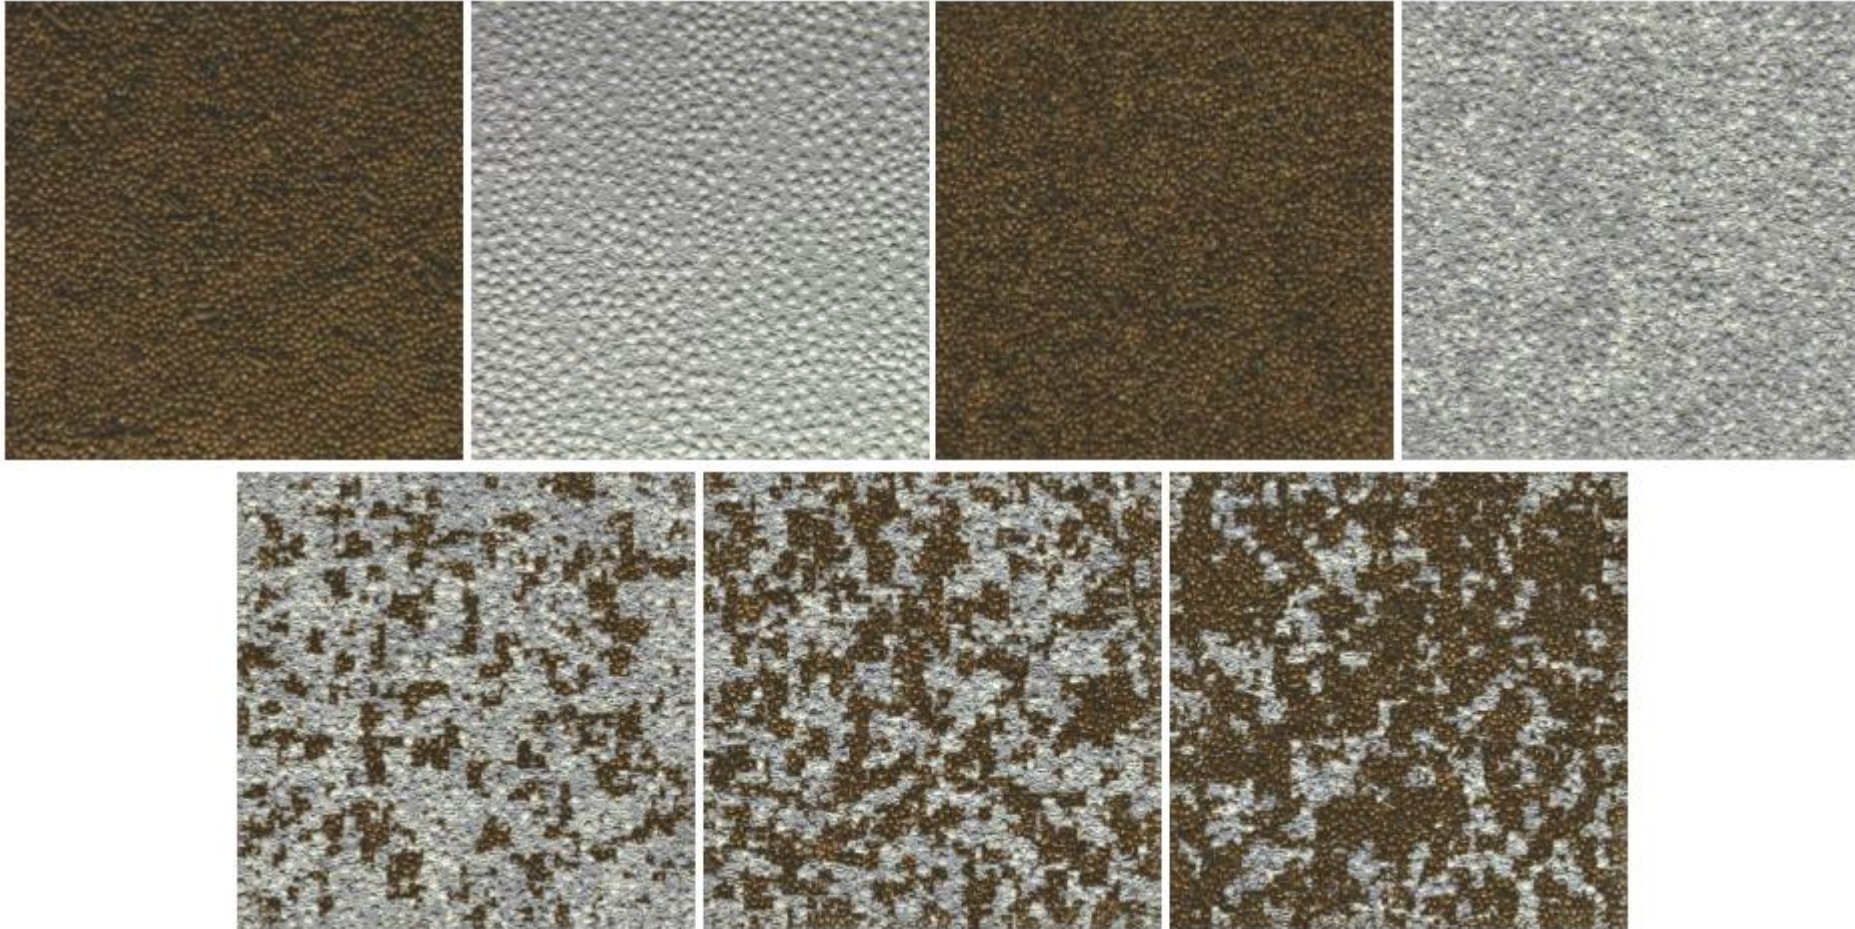


Figure 1: Source coding and simulation.

Learning from samples: Lempel-Ziv



Learning from samples: Lempel-Ziv



G. Brown, G. Sapiro, and G. Seroussi, "Texture Mixing via Universal Simulation," 2005

High Dimensions

Google Search results for "lav varshney". The search bar shows the query "lav varshney" and the URL is https://www.google.com/search?q=lav+varshney&rlz=1C1GCEA_enUS801US824&source=lnms&tbm=isch&sa=X&ved=2ahUKEwjWiduyq7v1AhXVHzQIHc46AFAQ...

Navigation tabs: All, News, Videos, Images, Maps, More. Tools: Collections, SafeSearch.

Filter tabs: twitter, bio, artificial intelligence, science festival, kush varshney, ibm, neural computing, uiuc, urbana champaign, illinois, coordinated science, illinc.

Search results (Images):

- Lav R Varshney | Electrical & Computer ...
ece.illinois.edu
- Lav Varshney (@lrvarshney) / Twitter
mobile.twitter.com
- Lav Varshney, PhD, Univer...
midas.umich.edu
- Lav Varshney – Center for Health...
chi.healtheng.illinois.edu
- Lav Varshney - IMDb
imdb.com
- Can AI help fix the food system? - T...
gfi.org
- STIR Member :: Lav Varshney
rle.mit.edu
- Lav Varshney (@varshney_lav) / Tw...
twitter.com

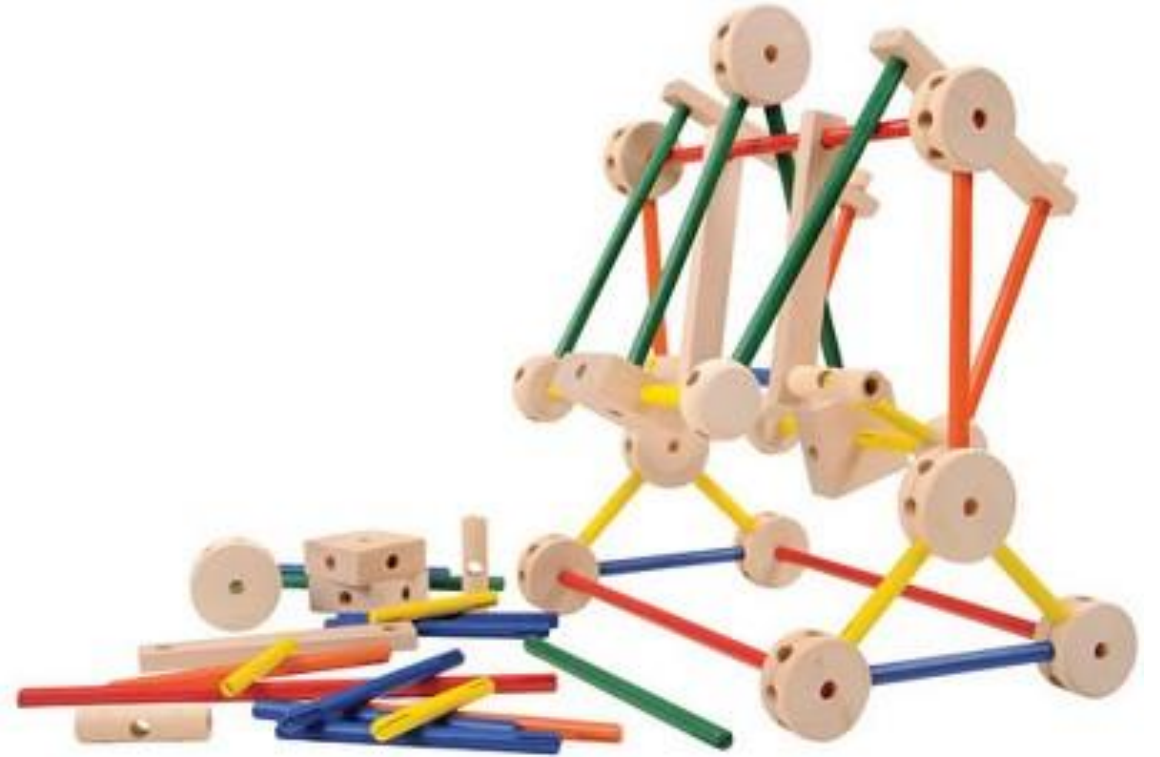
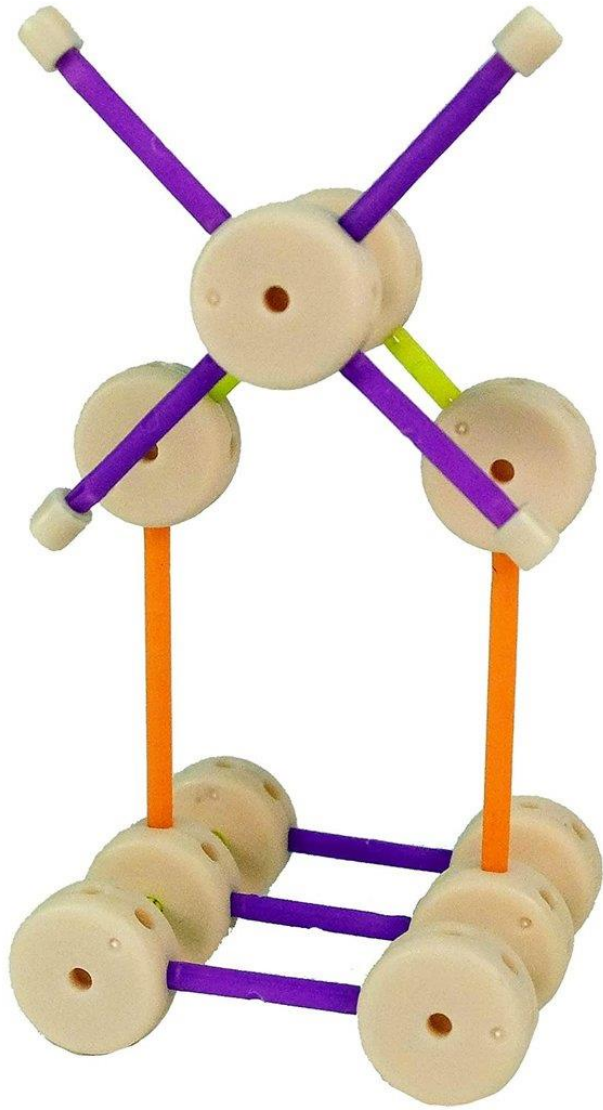
Search results (Videos):

- First Nerve: Twin Paradox: The Olfacto ...
firstnerve.com
- Can AI help fix the food system? - The ...
gfi.org
- Illinois ECE Professor Lav Varshney ...
youtube.com
- Lav Varshney (@lrvarsh...
mobile.twitter.com
- IBM 5 in 5 2012: Taste | IBM Rese...
ibm.com
- Lav Varshney (lvarsh) - ...
pinterest.com
- LIDS/ALL 2013
lidsmag.lids.mit.edu

Search results (Other):

- ICLR: BERTology Meets Biology: Interpreting Attention in Protein Language Models
BERTology Meets Biology: Interpreting Attention in Protein Language Models
Jesse Vig, Ali Madani, Lav R. Varshney, Caiming Xiong, Richard Socher, Harrison Falciani Rajani
- LAV VARSHNEY
- ILLINOIS INNOVATORS PODCAST WITH MIKE KOON
- The Videa to value Podcast
Host: Nick Vallicorn
Guest: Lav Varshney
#024 Prof Lav Varshney - How computers are becoming creative
- Lav Varshney (lvarsh) - ...
pinterest.com
- AI for Good
Global AI Summit

Structured (combinatorial) conceptual spaces



Extrapolation rather than interpolation

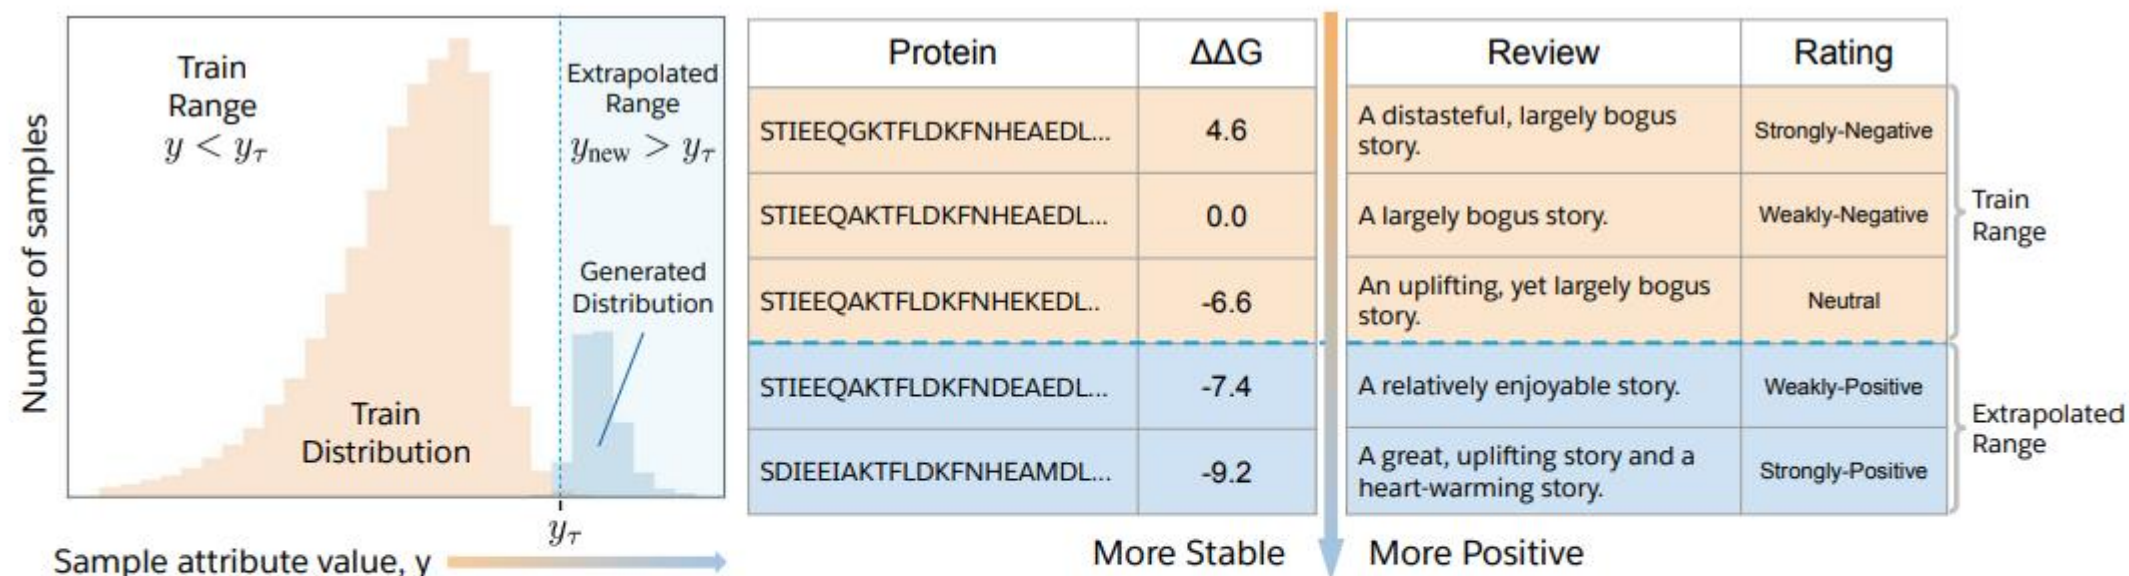


Figure 1: Attribute-enhanced generation. The goal of extrapolation (left) is to generate samples whose attribute values exceed that of all training samples, y_τ . We explore target attribute extrapolation for protein sequences (center) and movie reviews (right), where more stable protein sequences and more positives text reviews are generated.

What are generative AI models good for?

Why are you taking this elective class?



SYLLABUS

And Now, From I.B.M., Chef Watson



Robert Caplin for The New York Times

I.B.M. plans to serve a breakfast pastry devised by Watson and the chef James Briscione at its meeting on Thursday.

By STEVE LOHR

Published: February 27, 2013

I.B.M.'s Watson beat "Jeopardy" champions two years ago. But can it whip up something tasty in the kitchen?

More Tech Coverage

News from the technology industry, including start-ups, the Internet, enterprise and gadgets. On Twitter: @nytimesbits.



That is just one of the questions that I.B.M. is asking as it tries to expand its artificial intelligence technology and turn Watson into something that

actually makes commercial sense.

FACEBOOK

TWITTER

GOOGLE+

SAVE

E-MAIL

SHARE

PRINT

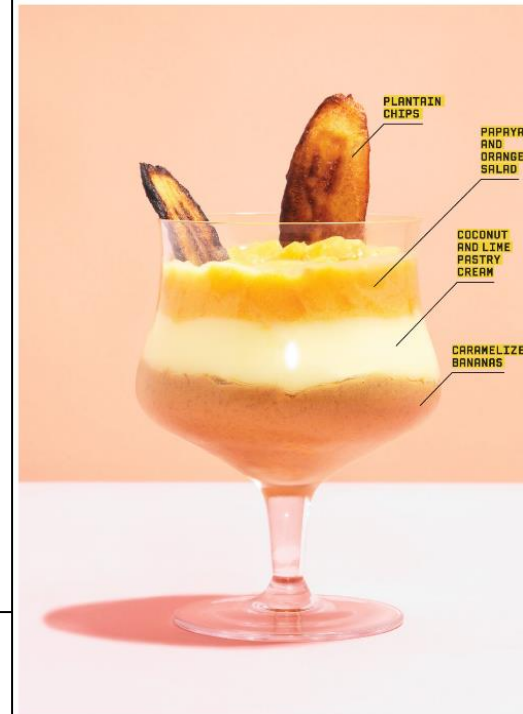
REPRINTS

SCIENCE | food | wired magazine

Share 486 Tweet 197 +1 30 LinkedIn Share 35 Pin it

Digital Gastronomy

WHEN AN IBM ALGORITHM COOKS, THINGS GET COMPLICATED—AND TASTY.



Prop styling: Laurie Raab | Justin Fantl

IBM's AI-like computer systems aren't limited to Watson, the *Jeopardy*-winning supercomputer that schooled Ken Jennings on national television. In fact, IBM researchers foresee a not-so-distant future when algorithms will be a replacement for inefficient customer service models, a diagnostic tool for doctors, and believe it or not, chefs.

Researcher Lav Varshney has already built an algorithm that creates recipes from parameters like cuisine type, dietary restrictions, and course. The system determines optimal mixtures based on three things: tens of thousands of recipes taken from sources like the Institute of Culinary Education or the Internet, a database of hedonic psychophysics (what humans like to eat), and food chemistry. Right now, the result is like a pre-Julia Child cookbook, providing chefs, who already know cooking basics, with suggestions for billions of ingredient combinations but no instructions.

To test its skill, we pitted IBM's algorithm against go-to-recipe resource Epicurious (owned by WIRED's parent company, Condé Nast). We searched the site for a Caribbean plantain dessert and found a tasty concoction with rum and coconut sauce. With the same parameters, IBM's computer generated a list of about 50 ingredients, including orange, papaya, and cayenne pepper, from which IBM researcher and professional chef Florian Pinel developed a mind-blowing Caymanian parfait. While the IBM dessert tasted better, it was also insanely elaborate, so we'll call it a draw.

—Allison P. Davis



IBM'S TASTE MASTER

COGNITIVE COMPUTING TAKES ON A NEW FRONTIER: MEAL PLANNING
BY VALERIE ROSS

[The New York Times, 27 Feb. 2013]
[San Jose Mercury News, 28 Feb. 2013]
[IEEE Spectrum, 31 May 2013]
[Wired, 1 Oct. 2013]







Cognitive Cooking with Chef Watson

Recipes for Innovation from IBM & the Institute of Culinary Education



look for
INGREDIENTS

choose a
DISH

pick a
STYLE

start
AGAIN

FAVORITES

STRAWBERRY

Based on: pad thai noodles from Bon Appétit

Strawberry Curry


HERE'S A STARTING POINT * . . .

8 servings

SWEETENER 1¼ tsp acacia honey	FISH/SEAFOOD 2½ lb king salmon	OIL/FAT 1 tbsp peanut oil
FRUIT strawberry	HERB chive	STOCK/SOUP ½ cup chicken broth
DAIRY 1½ cup yogurt	VEGETABLE 1 diced, pitted, peeled avocado 22 sliced, trimmed radishes ½ cucumber (8-1¼"), chopped english cucumber	CONDIMENT 2½ tbsp curry paste 1½ oz red curry paste

ALCOHOLIC BEVERAGE
bourbon

1. Chef Watson is pretty sure that bourbon will taste good in this dish, but needs your help in figuring out the details.
2. Heat peanut oil in heavy large skillet over medium-high heat.
3. Add sliced radishes; stir-fry.
4. Reduce heat to medium.
5. Add english cucumber and red curry paste; stir about 1 minute.



Based On Thai Shrimp Curry
From Bon Appétit

TUTORIAL

PRIVACY

TERMS OF USE

MORE...

MORE...

MORE...

MORE...

<https://www.ibmchefwatson.com>

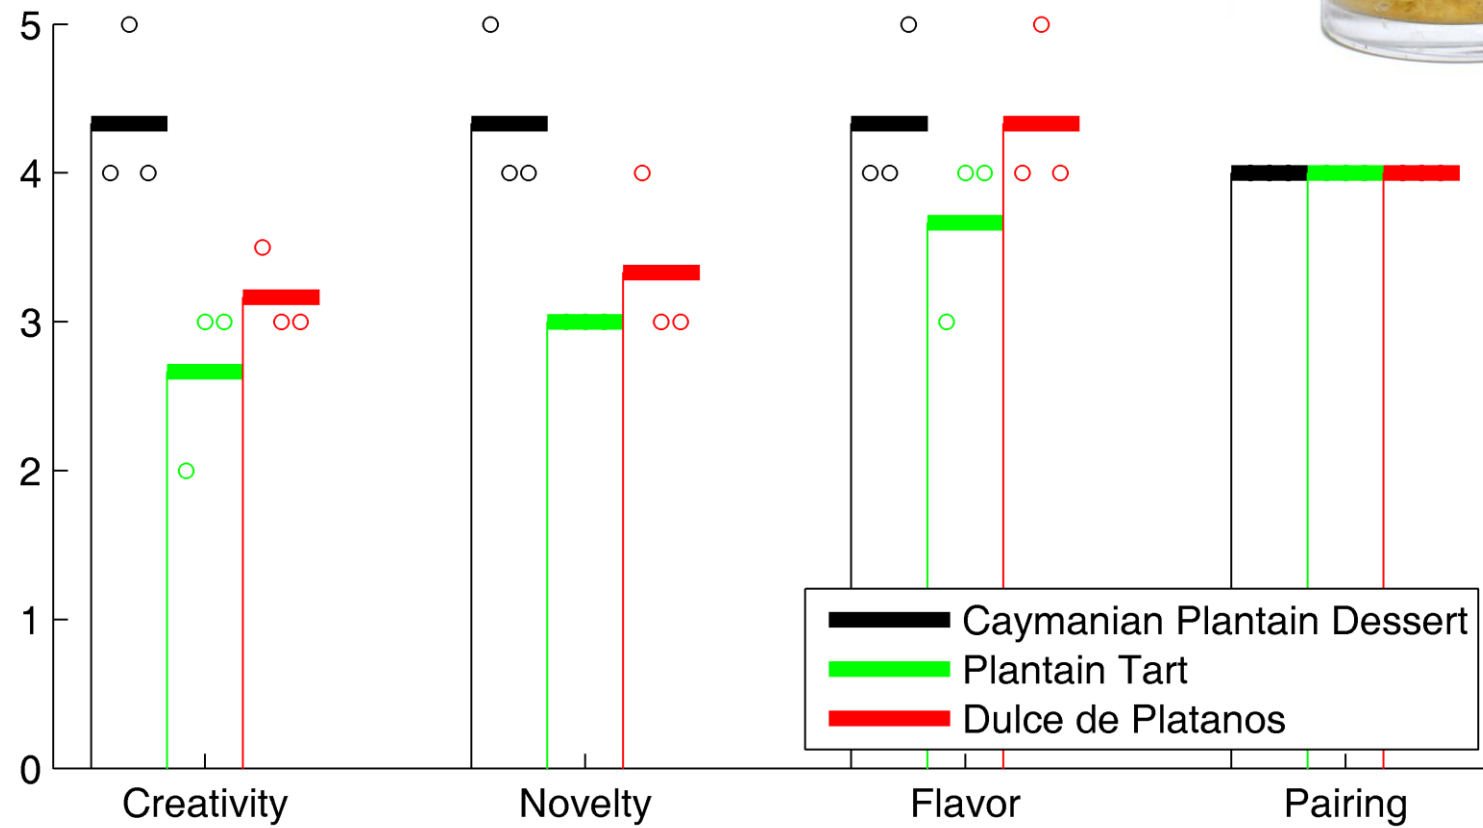


prepared with
IBM Chef Watson™





Consensual assessment technique



Beyond the Turing Test: Lovelace 2.0



LOVELACE

Lovelace: “only when computers originate things should they be believed to have minds”



FERRUCCI

Lovelace 1.0: an artificial agent possesses intelligence in terms of whether it can “take us by surprise”



RIEDL

Lovelace 2.0: An artificial agent must create artifact o of type t where:

- artifact o conforms to constraints C where $c_i \in C$ is any criterion expressible in natural language
- human evaluator h , having chosen t and C , is satisfied o is valid instance of t and meets C , and
- human referee r determines combination of t and C to not be impossible



BON APPÉTIT / ENTERTAINING + STYLE / TRENDS + NEWS

5:59 AM / JUNE 30, 2014



Former IBM Research scientist Lav Varshney presents a demo of an early version of the cognitive cooking technology at IBM Research.

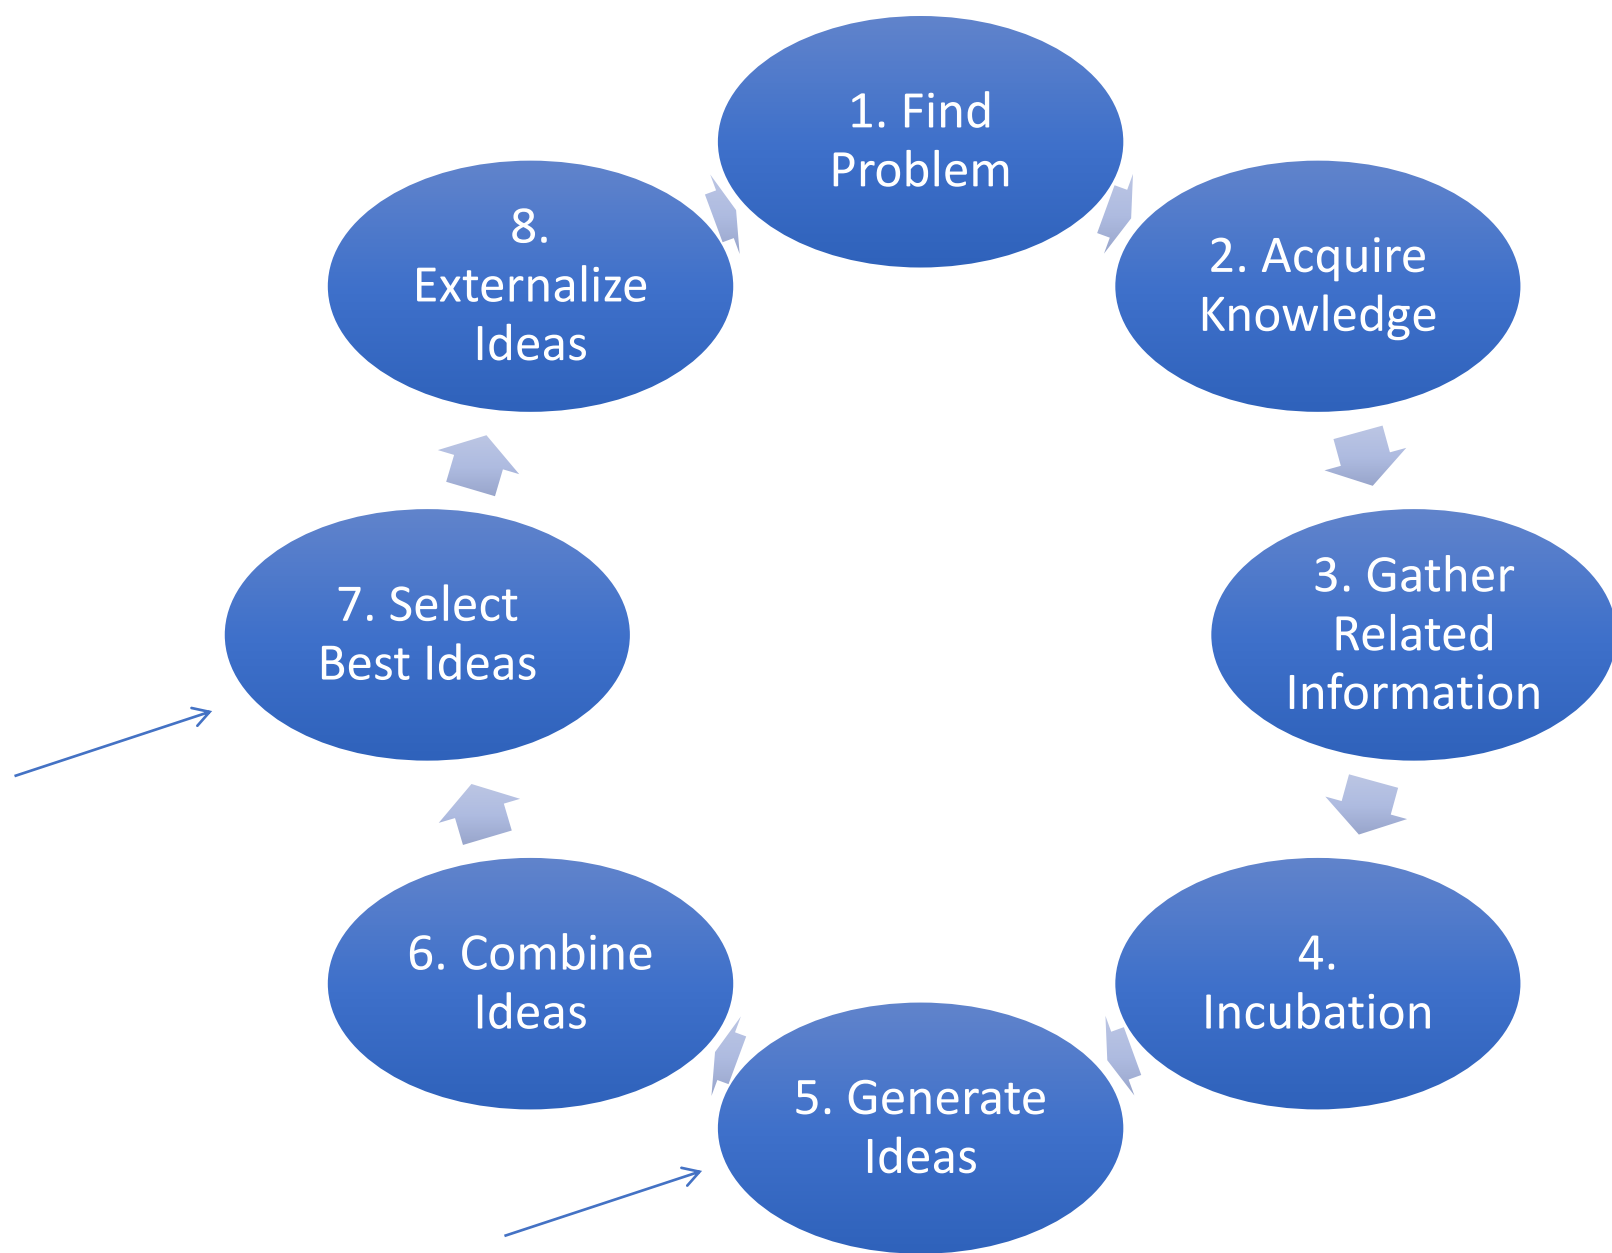
CREDIT: COURTESY IBM

How IBM's Chef Watson Actually Works

[f](#) [t](#) [p](#) [g+](#) [x](#) [+](#) 47 [Comment](#)

WRITTEN BY ROCHELLE BILOW

[PRINT](#) [RSS](#)



[Sawyer, 2012]



BON APPÉTIT / ENTERTAINING + STYLE / TRENDS + NEWS

5:59 AM / JUNE 30, 2014



Former IBM Research scientist Lav Varshney presents a demo of an early version of the cognitive cooking technology at IBM Research.

CREDIT: COURTESY IBM

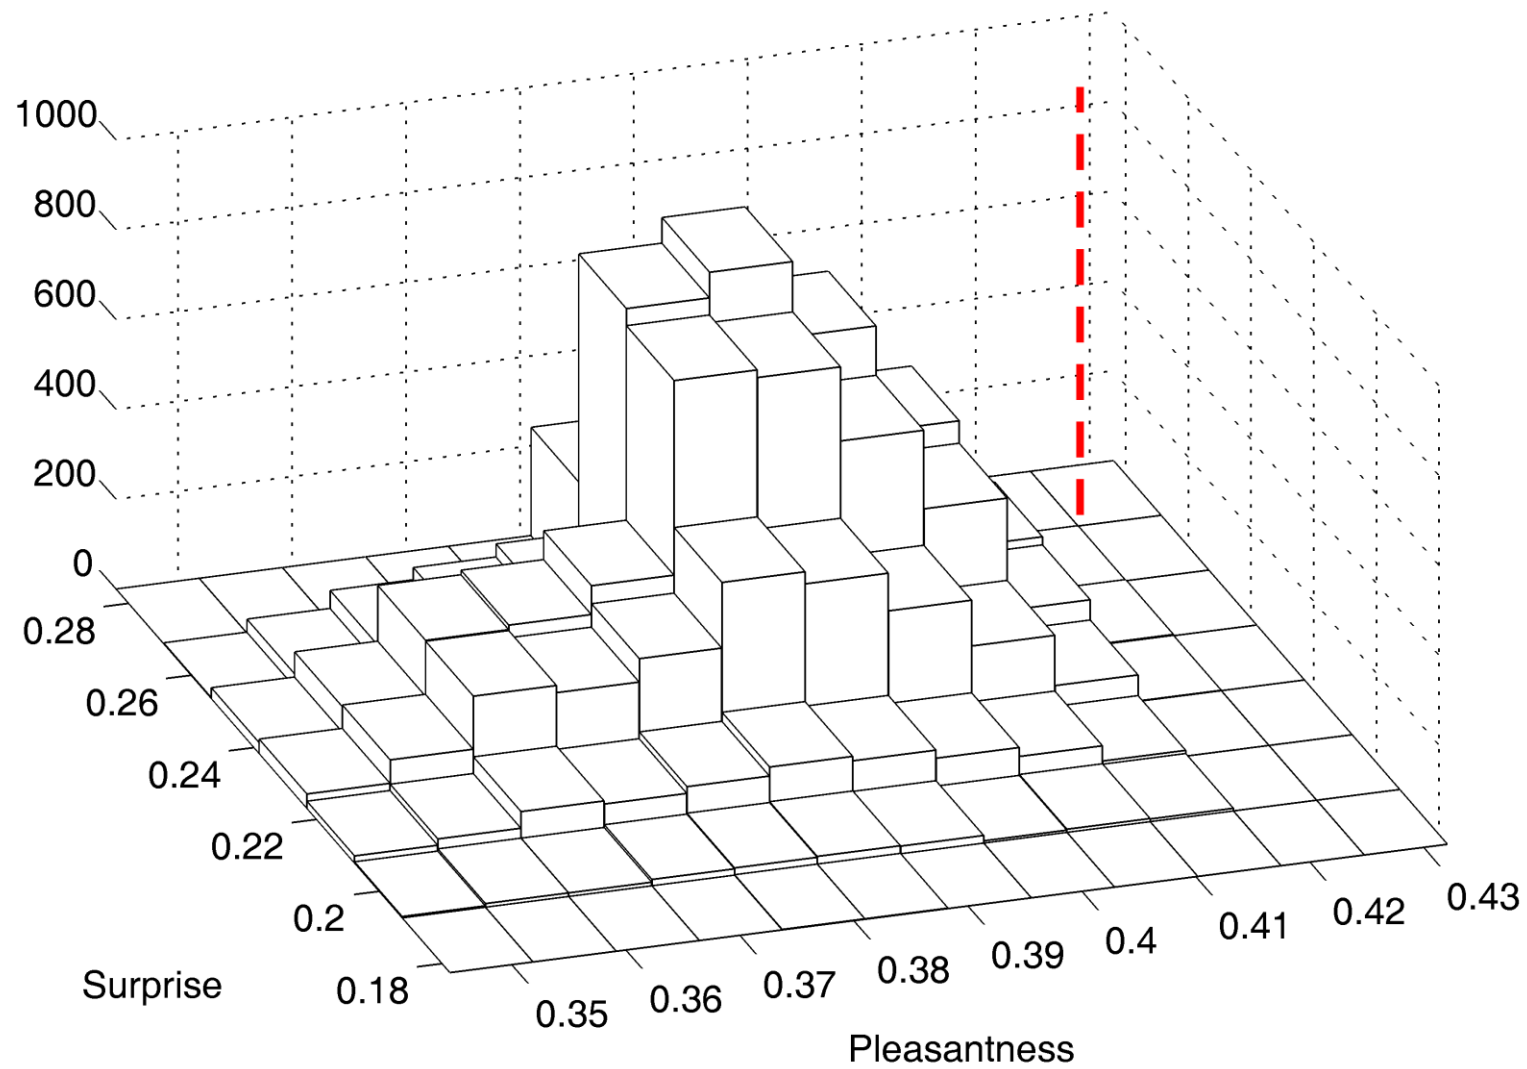
How IBM's Chef Watson Actually Works



WRITTEN BY ROCHELLE BILOW

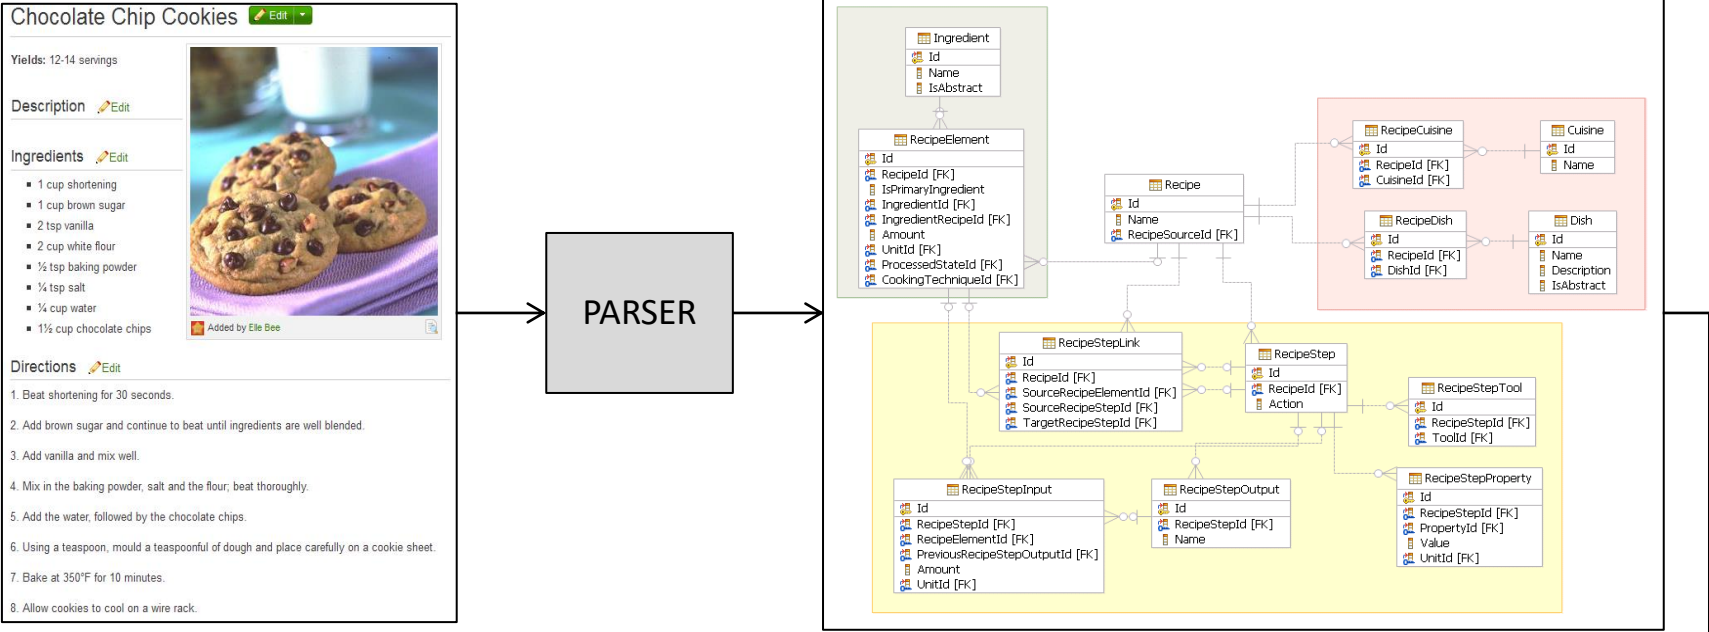
[PRINT](#) [RSS](#)

1. Sample from state space, using culturally well-chosen sampling distribution
2. Rank according to psychophysical predictors of novelty and flavor
3. Select either automatically or semi-automatically depending on human-computer interaction model

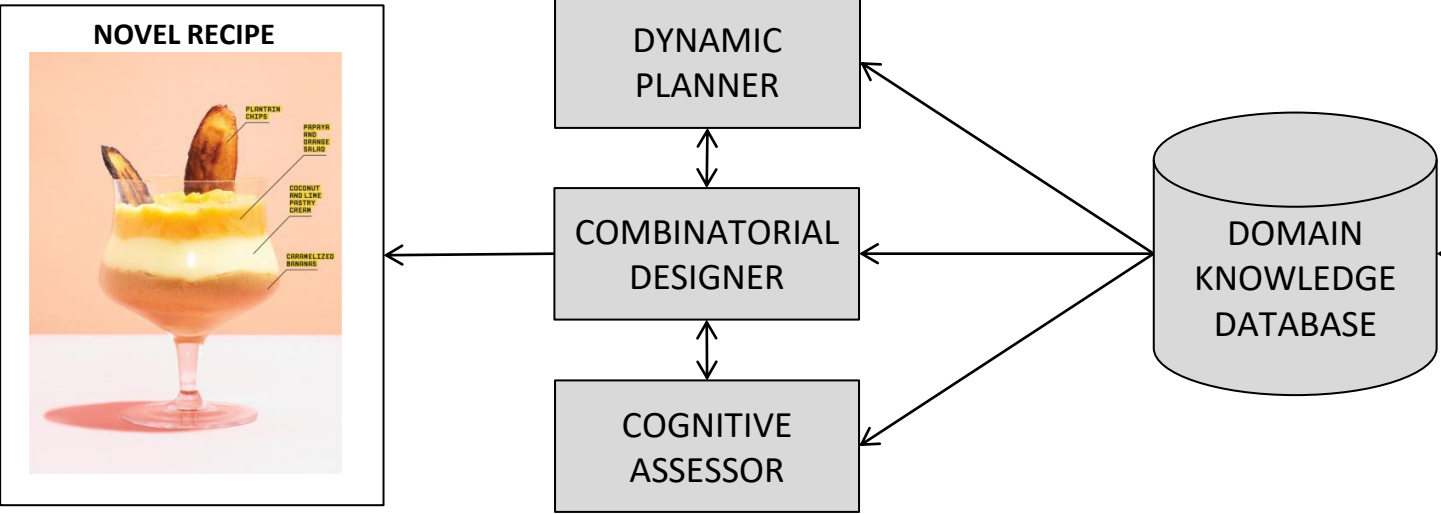


Joint histogram of surprise and pleasantness for 10000 generated Caymanian Plantain Dessert recipes. Values for the selected/tested recipe indicated with red dashed line.

Data Engineering and Natural Language Processing to Understand the Domain



Generative, Selective, and Planning Algorithms to Create the Best New Ideas



KEY INGREDIENT: ROOT VEGETABLES

FRIED LOTUS ROOT CHIPS

Yield: Makes a lot

2 lotus roots, peeled

Vegetable oil to fry

Kosher salt to taste

Pinch of cayenne pepper, optional

Thinly slice the lotus root using a mandolin. (If not frying right away, hold the lotus root in water with some vinegar or lemon juice to prevent oxidation.) Heat 2 inches of oil in a heavy pot to 360° F. Pat the sliced lotus root dry with paper towel, and fry in batches until golden brown (they will continue to brown once removed, so cook just to golden). Transfer to a rack over a rimmed sheet pan, and sprinkle with salt (mix in a bit of cayenne pepper to the salt, if a spicier chip is desired).

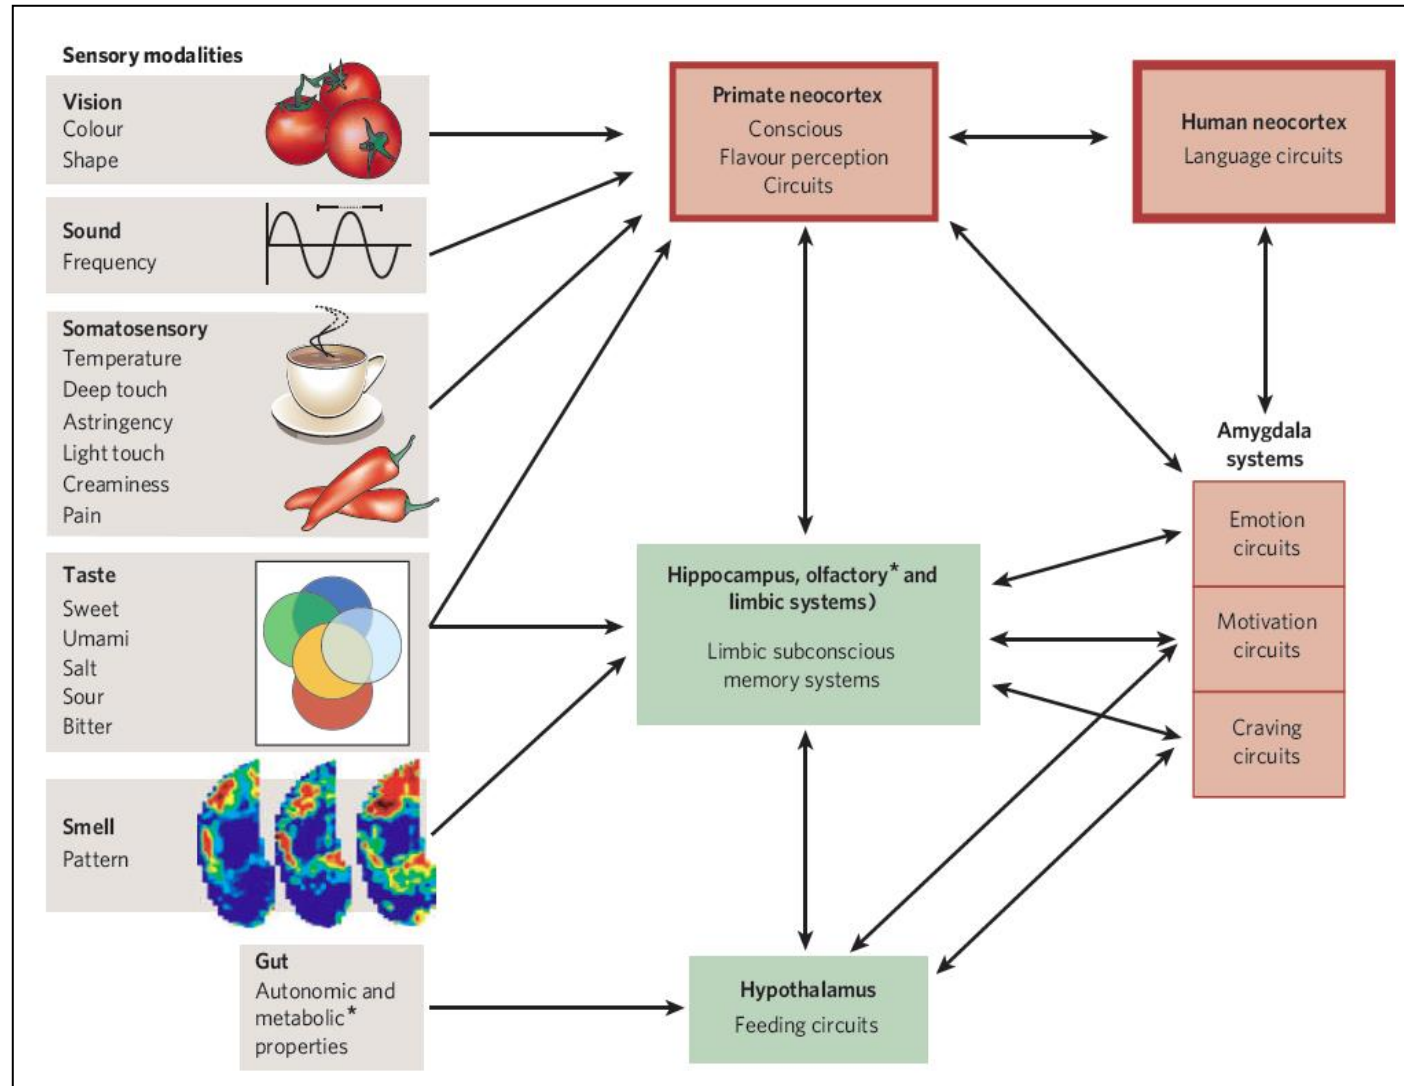
Institute

50

New Yo

www

Neurogastronomy

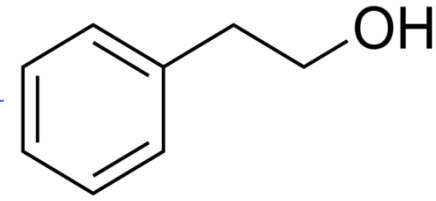


[Shepherd,
2006]

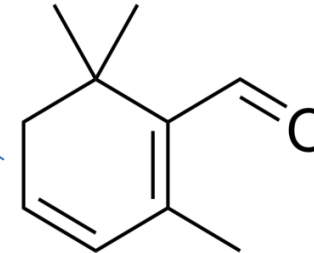
Food Chemistry

Saffron (*Crocus sativus* L.)

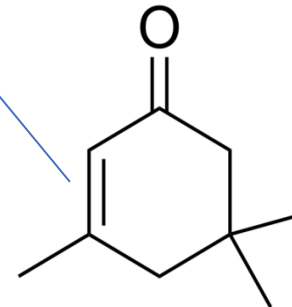
2-phenylethanol (=phenethyl alcohol)
safranal (=2,6,6-trimethyl-1,3-cyclohexadienecarbaldehyde)
3,5,5-trimethyl-2-cyclohexen-1-one (=isophorone)
hexadecanoic acid (=palmitic acid)
2,6,6-trimethyl-2-cyclohexene-1,4-dione
(Z,Z)-9,12-octadecadienoic acid (=linoleic acid)
(Z,Z,Z)-9,12,15-octadecatrienoic acid (=linolenic acid)
naphthalene
2,4,6-trimethylbenzaldehyde (=mesitylaldehyde)
2,6,6-trimethyl-1,4-cyclohexadienecarbaldehyde
6,6-dimethyl-2-methylene-3-cyclohexenecarbaldehyde
4-hydroxy-2,6,6-trimethyl-1-cyclohexenecarbaldehyde (=4-hydroxysafranal)
3,5,5-trimethyl-3-cyclohexen-1-one
3,3,4,5-tetramethylcyclohexanone
3,5,5-trimethyl-4-methylene-2-cyclohexen-1-one
4-hydroxy-3,5,5-trimethyl-2-cyclohexen-1-one
2,3-epoxy-4-(hydroxymethylene)-3,5,5-trimethylcyclohexanone
5,5-dimethyl-2-cyclohexene-1,4-dione
2,2,6-trimethylcyclohexane-1,4-dione (=3,5,5-trimethyl-cyclohexane-1,4-dione)
2-hydroxy-3,5,5-trimethyl-2-cyclohexene-1,4-dione
2-hydroxy-4,4,6-trimethyl-2,5-cyclohexadien-1-one
2,6,6-trimethyl-3-oxo-1,4-cyclohexadienecarbaldehyde
4-hydroxy-2,6,6-trimethyl-3-oxo-1,4-cyclohexadienecarbaldehyde
4-hydroxy-2,6,6-trimethyl-3-oxo-1-cyclohexenecarbaldehyde
3-hydroxy-2,6,6-trimethyl-4-oxo-2-cyclohexenecarbaldehyde
4-(2,2,6-trimethyl-1-cyclohexyl)-3-buten-2-one
4-(2,6,6-trimethyl-1-cyclohexen-1-yl)-3-buten-2-one (=β-ionone)
verbenone (=2-pinen-4-one)
octadecanoic acid (=stearic acid)
(Z)-9-octadecenoic acid (=oleic acid)
2(5H)-furanone (=crotonolactone, 2-buten-4-olide, 4-hydroxy-2-butenic acid lactone)



phenethyl alcohol

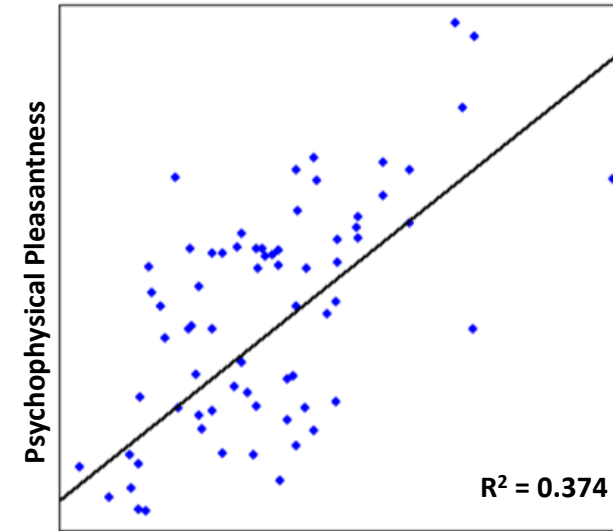
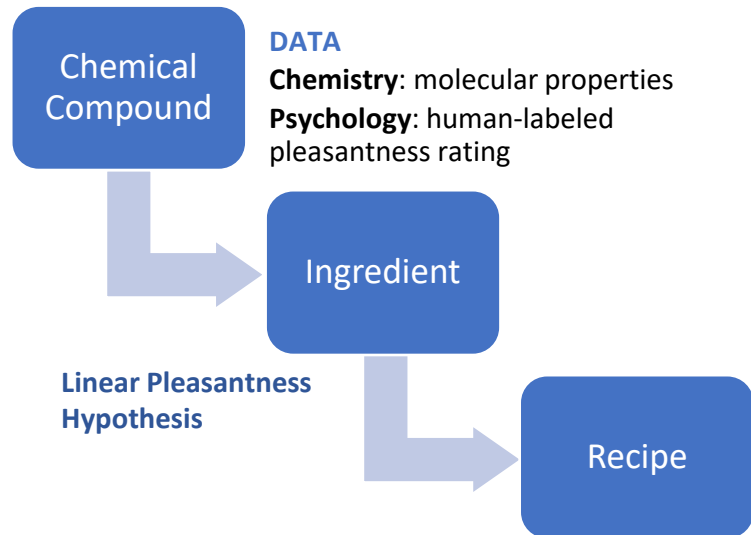
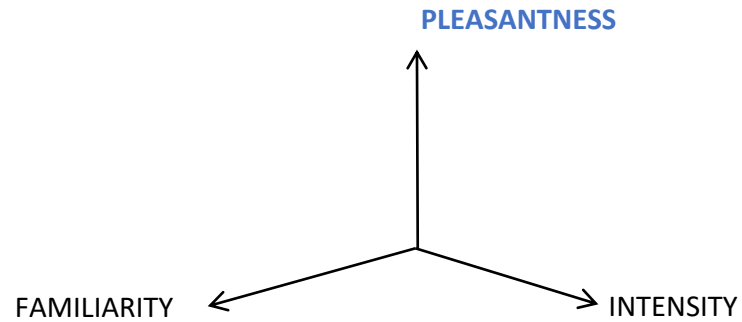


safranal



isophorone

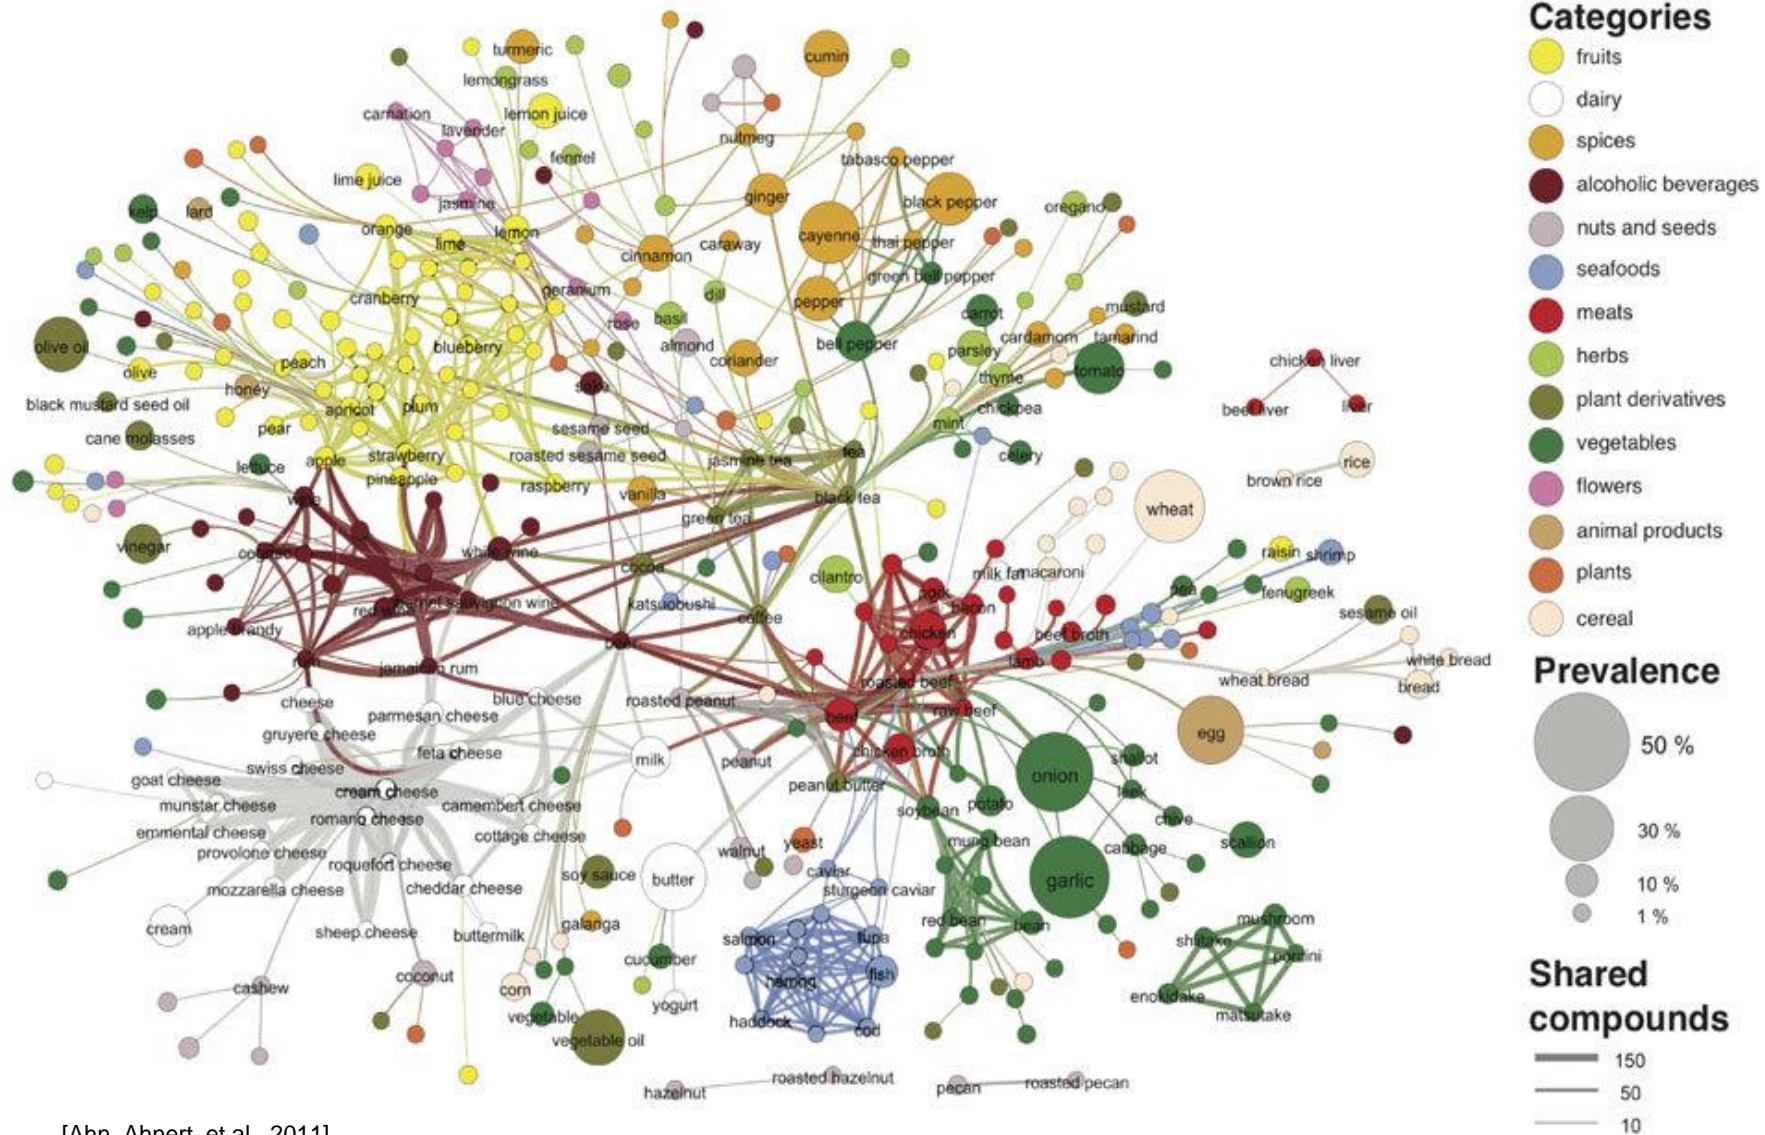
Hedonic Psychophysics



Chemistry [TPSA, heavy atom count, complexity, rotatable bond count, hydrogen bond acceptor count]

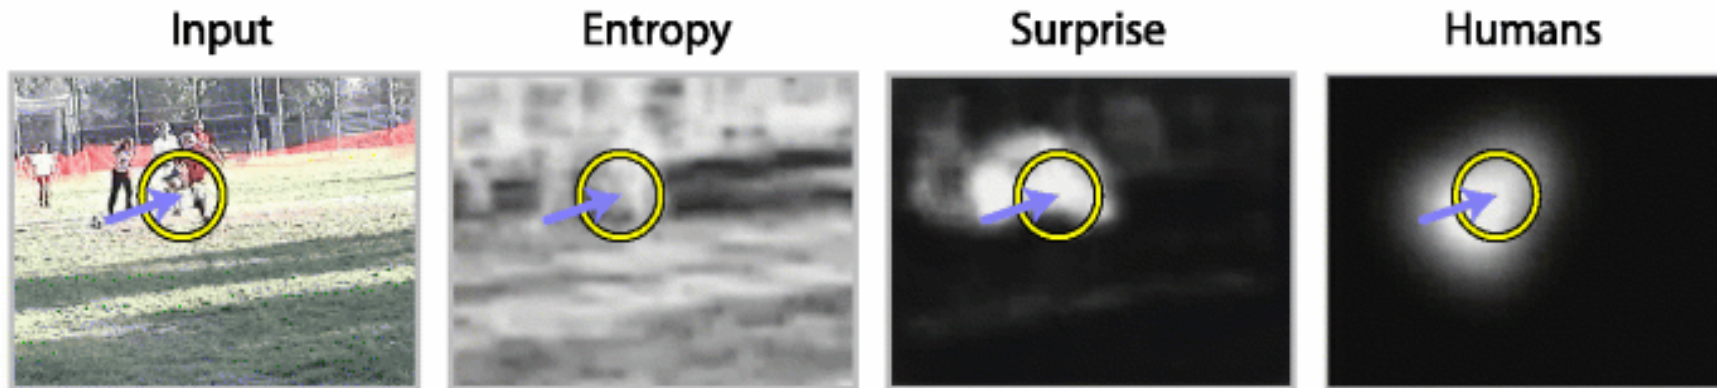
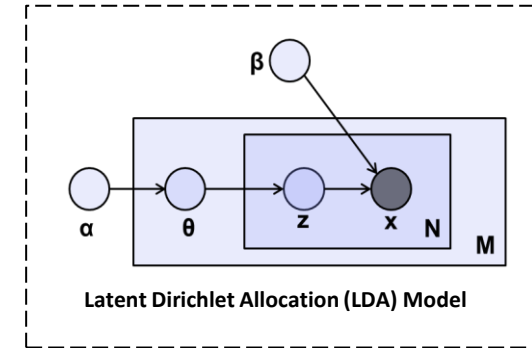
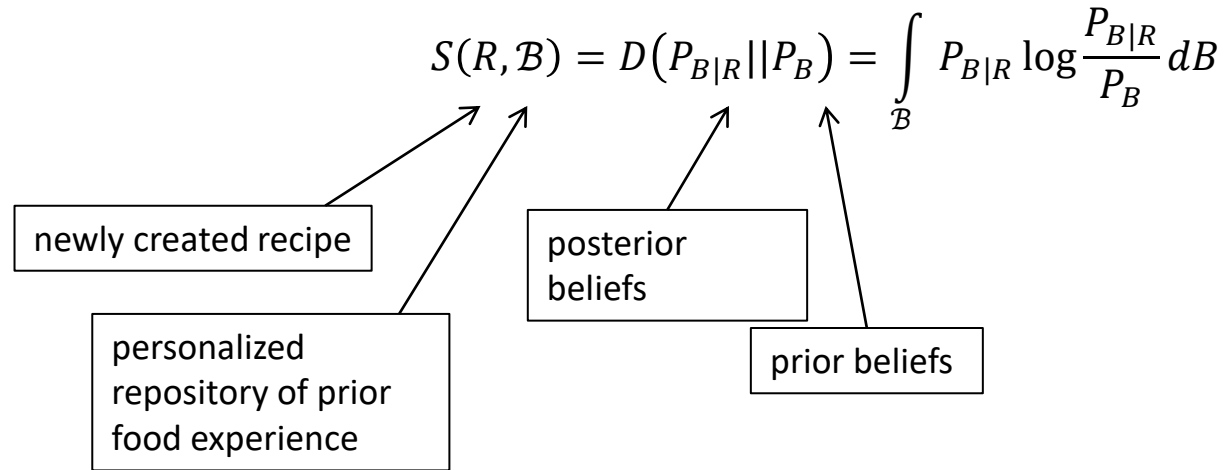
Black Tea
Bantu Beer
Beer
Strawberry
White Wine
Cooked Apple

Flavor Networks

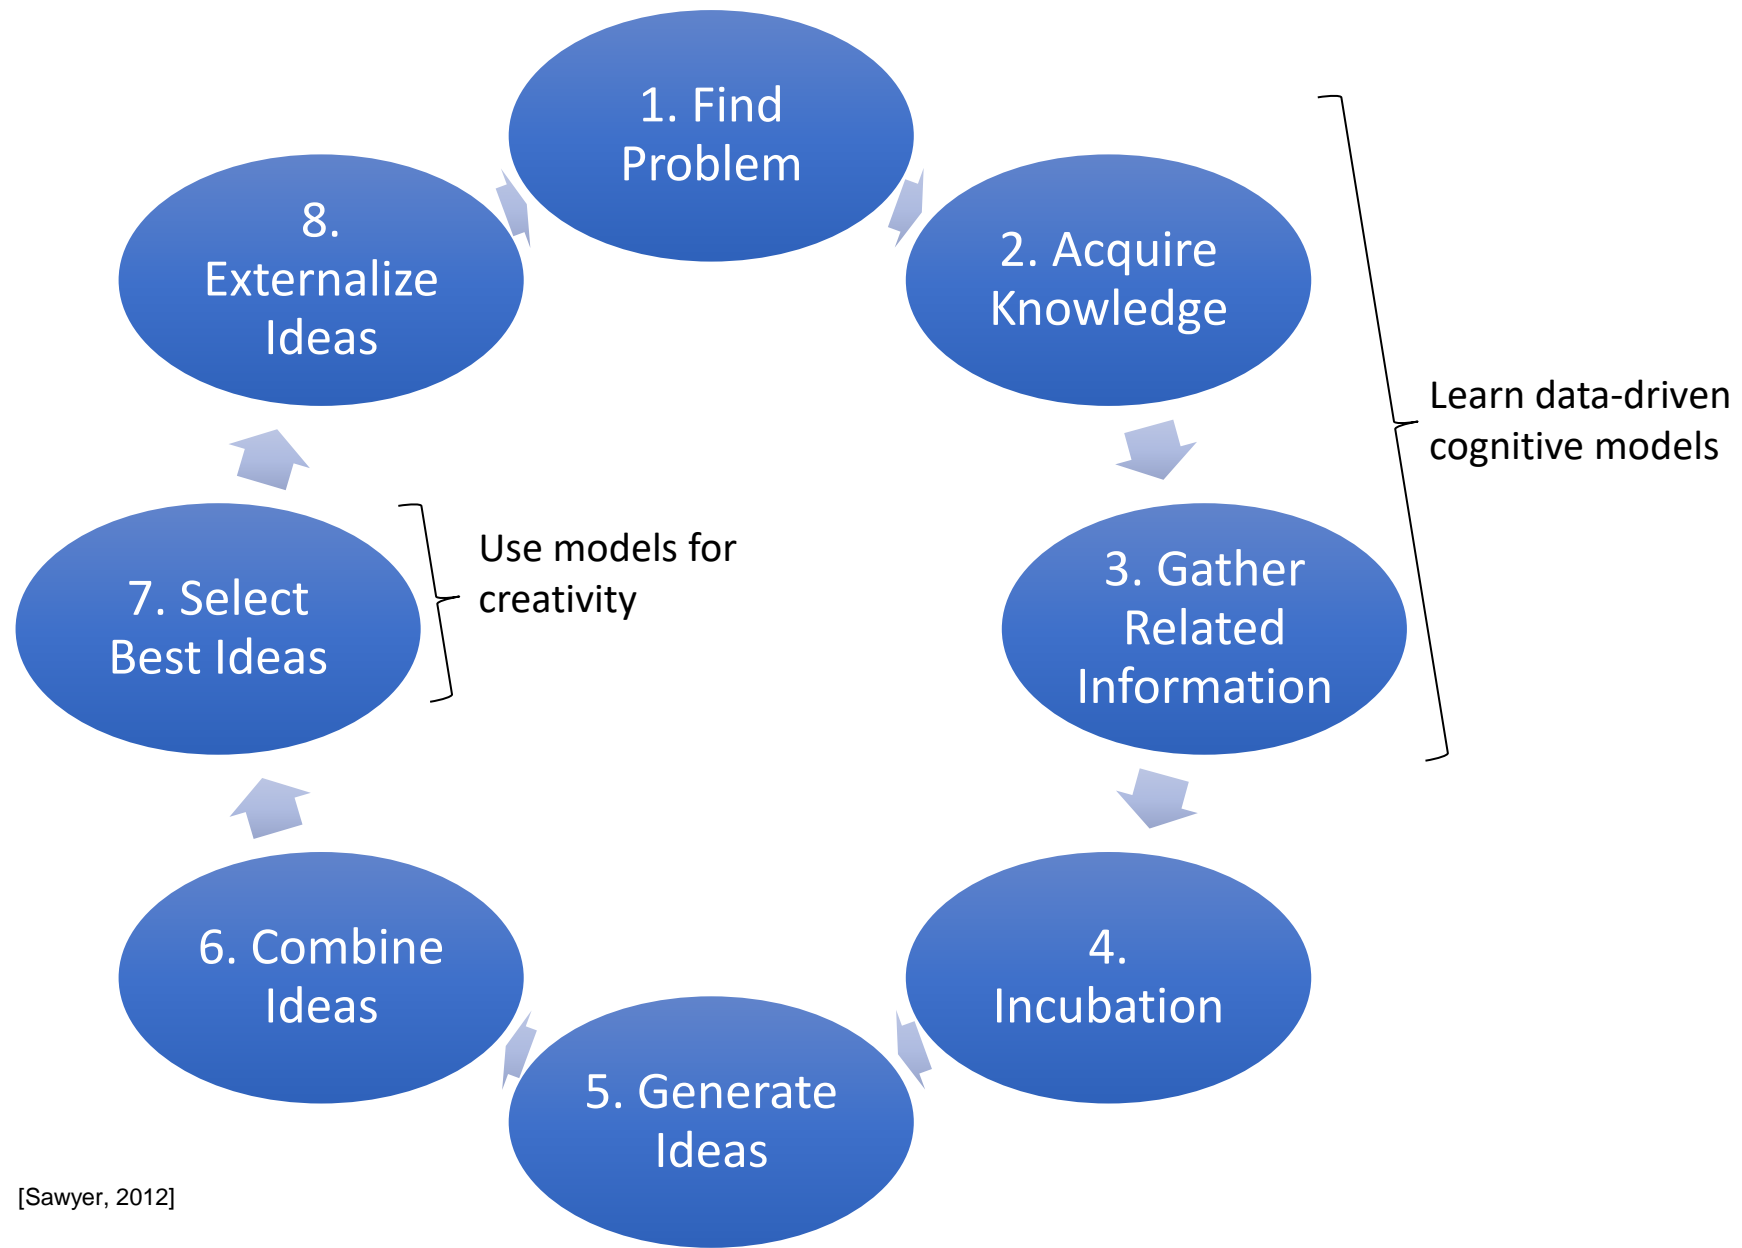


[Ahn, Ahnert, et al., 2011]

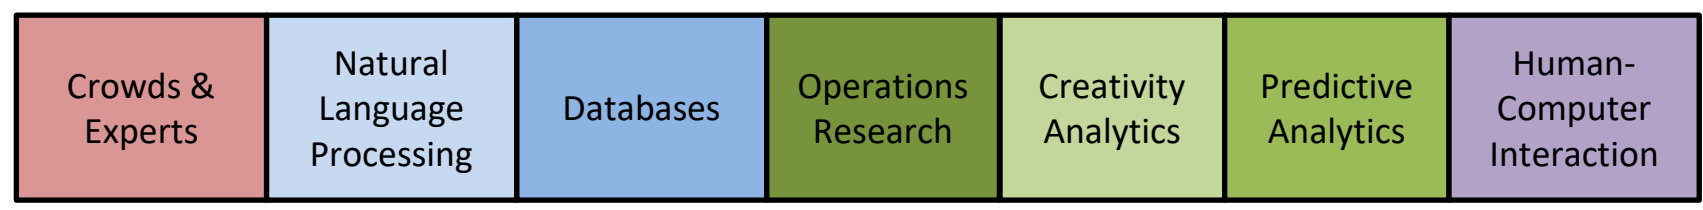
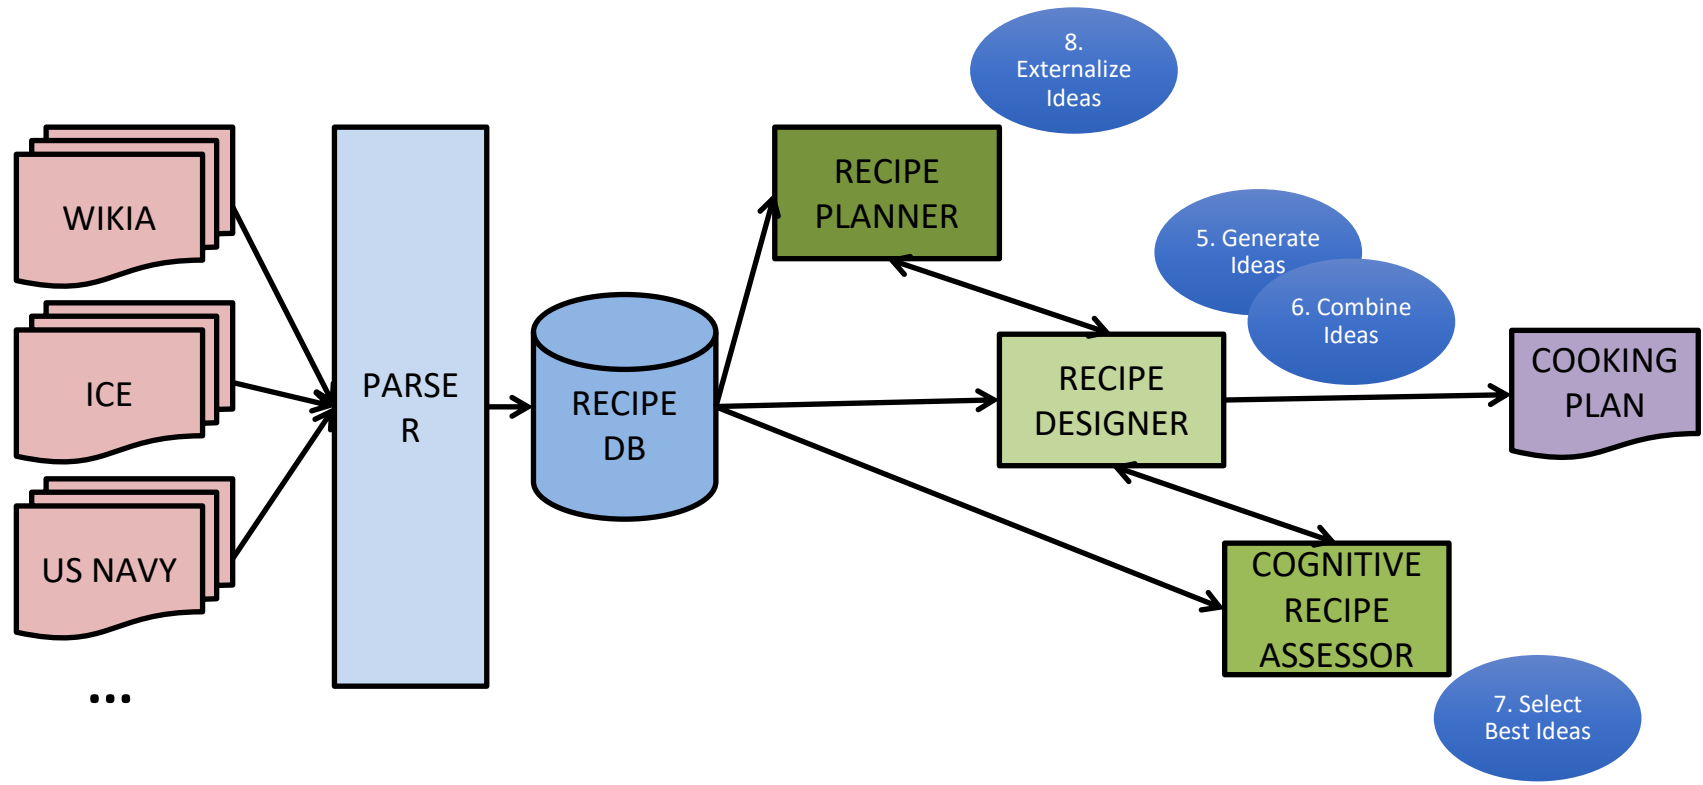
Bayesian Surprise and Attention



[Itti and Baldi, 2006]



[Sawyer, 2012]



Computational creativity / accelerated discovery for engineering materials

- At least 8% of global emissions caused by humans come from the cement industry alone
- Interest in reducing environmental impacts of construction materials while complying with product specifications.

The international journal of science / 30 September 2021

nature

Concrete needs to lose its colossal carbon footprint

Concrete will be crucial for climate-resilient construction. But the cement industry must set out its plan for decarbonization.

Wet concrete has been poured into buildings, roads, bridges and more for centuries. Structures using concrete have survived wars and natural disasters, outlasting many of the civilizations that built them¹. Alongside its strength and resilience, concrete is also a staple of building because it is relatively cheap and simple to make. Worldwide, 30 billion tonnes of concrete is used each year. On a per capita basis, that is 3 times as much as 40 years ago – and the demand for concrete is growing more steeply than that for steel or wood².

Versatile and long-lasting, concrete buildings and structures are in many ways ideal for climate-resilient construction. But concrete has a colossal carbon footprint – at least 8% of global emissions caused by humans come from the cement industry alone³. We must decarbonize its production.

Concrete is made by adding sand and gravel to cement,

cement production itself needs to be decarbonized, which could happen in a number of ways. For example, low-carbon fuels – such as hydrogen or biomass – could be substituted for fossil fuels in heating the limestone and clay. And scientists are examining whether electricity – instead of combustion – could be used for the heating.

Carbon capture could be part of the cement industry's transition process⁴. In Sweden, for example, a company announced in July that it wants to capture 1.8 million tonnes of CO₂ from a cement plant and bury it in the North Sea. Another possibility is to pump the captured CO₂ into concrete itself, locking it up forever – which might also improve the properties of the resulting material. The injected CO₂ reacts with calcium ions in the cement, producing more calcium carbonate, and potentially making the concrete able to withstand larger loads.

Concrete options

Technological changes can be accelerated through regulation and legislation. A huge proportion of concrete is used in public building projects. In North America alone, public agencies buy as much as one-third of concrete manufactured annually. That means they have leverage in the low-carbon transition: they could work with researchers and manufacturers to reshape the concrete industry.

In New York and New Jersey, a bill is making its way through state legislatures that, if passed, will mandate that state agencies and departments prioritize cement that has a lower carbon footprint.

Other states are introducing legislation that requires construction proposals to declare the environmental impact of cement mixes. Some regions, such as Honolulu,

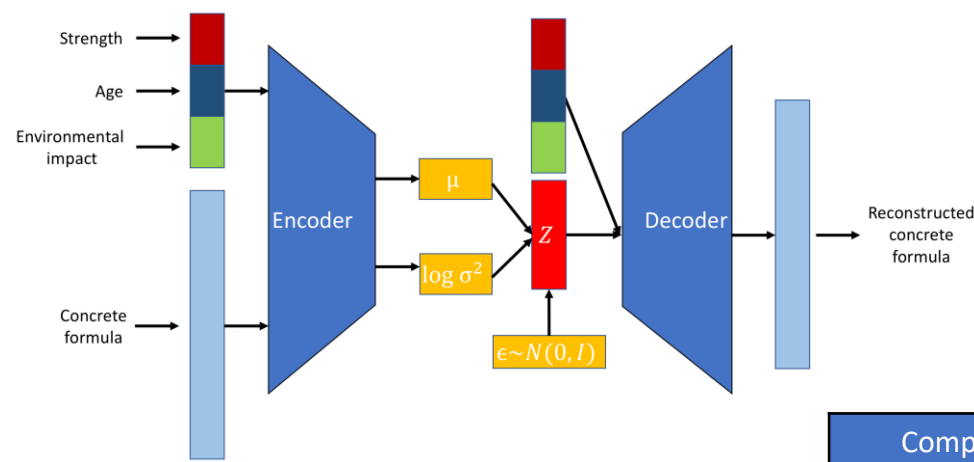
“At least 8% of global emissions caused by humans come from the cement industry alone.”

Small historical data set

- UCI ML repository concrete strength dataset + Environmental impact evaluated using the Cement Sustainability Initiative's Environmental Product Declaration tool:
 - 1030 instances
 - 8 input variables (composition)
 - 1 (compressive strength) + 12 (environmental impact) output variables

<https://github.com/IBM/Conditional-Variational-Autoencoder-for-Concrete-Design>

Conditional variational autoencoder (CVAE)

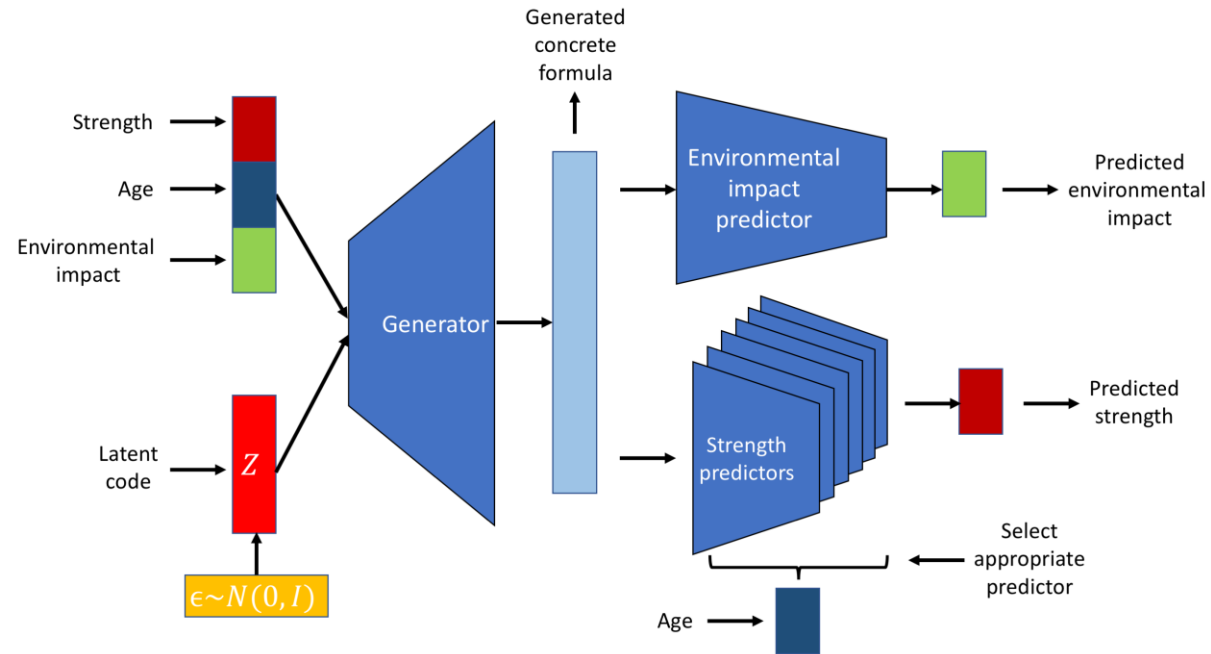


Strength	$[0,1]$
Age	$\{0,1\}^6$
Environmental Impact	$[0,1]^{12}$
Concrete formula	$[0,1]^7$

Want conditioning variables and extrapolative ability

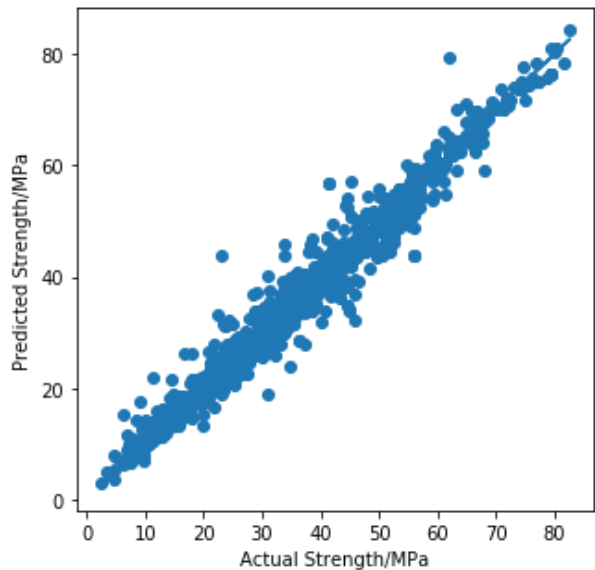
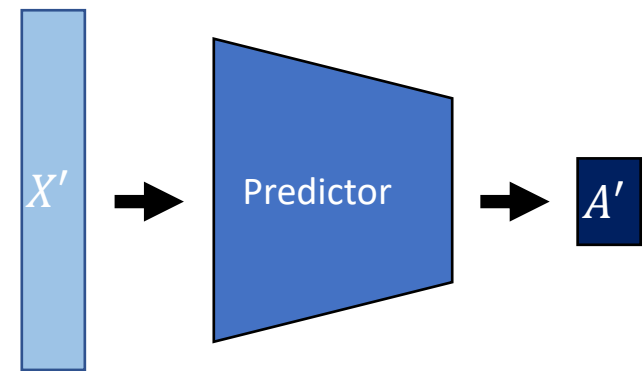
Component		Neuron type and activation	Dimension
Encoder	Layer 1	Fully connected+ReLU	26×25
	Layer 2	Fully connected+ReLU	25×20
μ		Fully connected+Softplus	25×2
$\log \sigma^2$		Fully connected+Softplus	25×2
Z		Reparameterization Trick	2
Decoder	Layer 1	Fully connected+ReLU	2×20
	Layer 2	Fully connected+ReLU	20×25
Output layer		Fully connected+Sigmoid	25×7

Property predictors



	Neuron type and activation	Dimension
Layer 1	Fully connected+ReLU	7×90
Layer 2	Fully connected+ReLU	90×35
Layer 3	Fully connected+ReLU	35×10
Output layer	Fully connected+ReLU	10×7

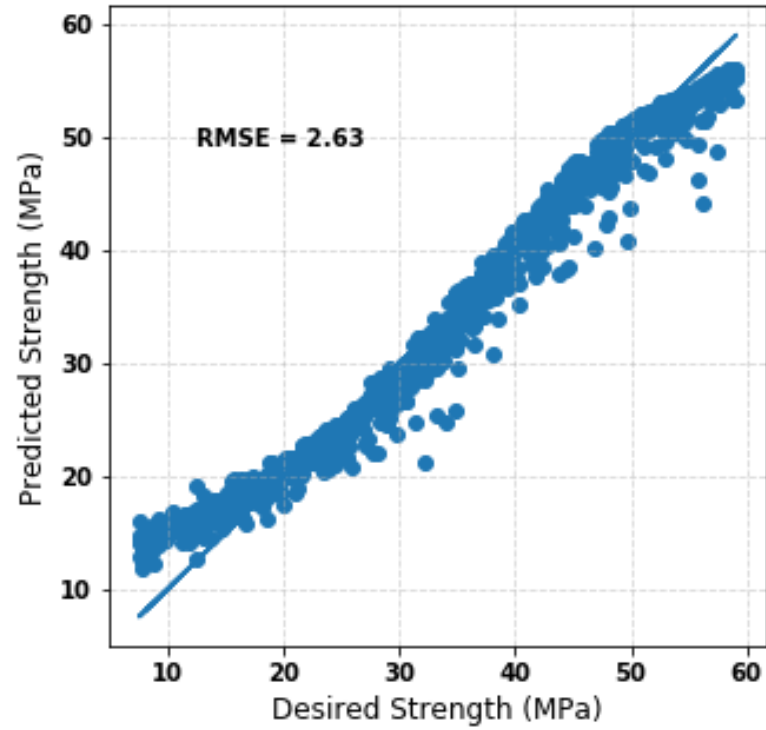
Strength regressor performance



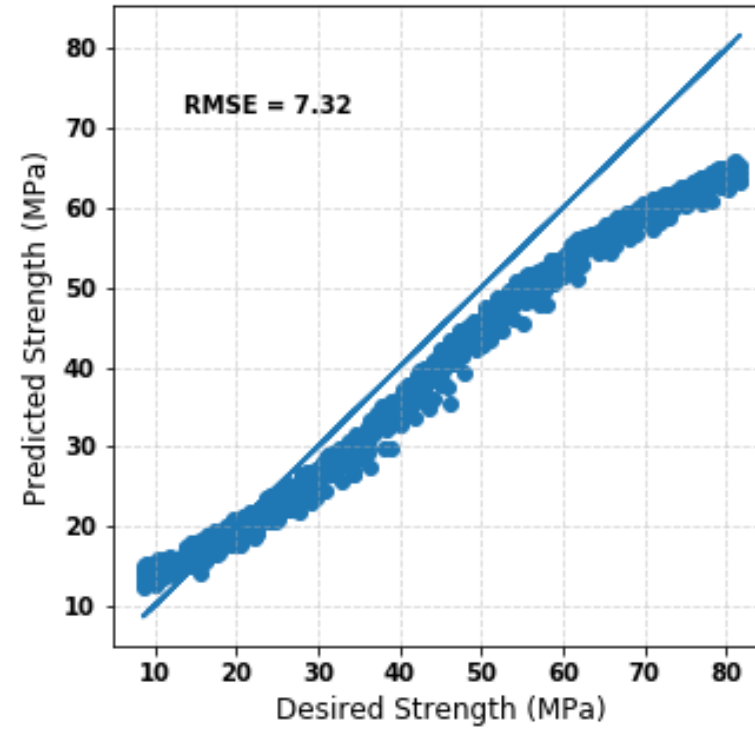
Metric	GWP (kg CO ₂ eq.)	AP (kg SO ₂ eq.)	CBW (m ³)
MAE	7.187	0.019	0.003
RMSE	9.374	0.040	0.006
R ²	0.979	0.974	0.881

	Predictor Performance (MPa)					
Metric	≤3	7	14	28	56	≥90
MAE	2.985	3.850	3.378	6.015	5.093	4.457
RMSE	0.222	0.201	0.163	0.227	0.124	0.125
R ²	0.819	0.870	0.703	0.679	0.795	0.789

Strength-conditioned progression

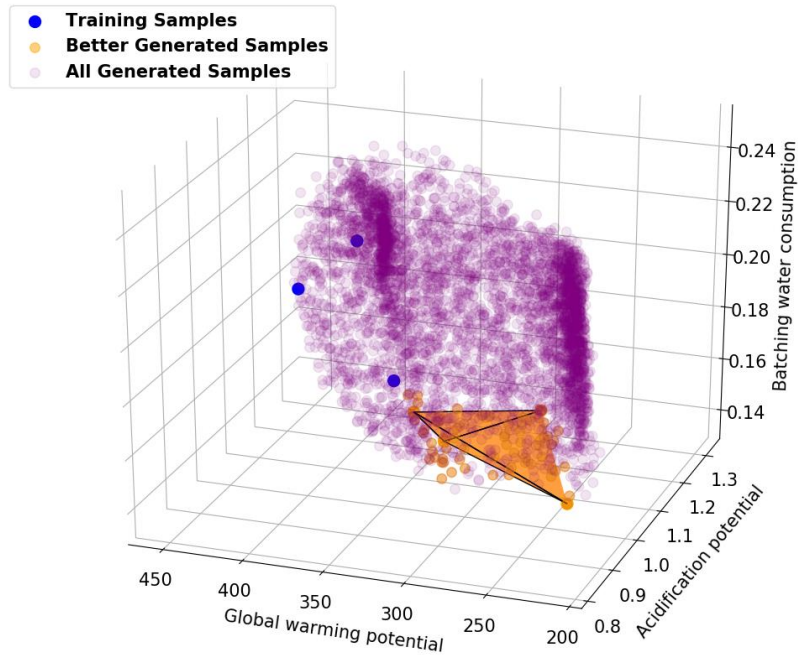


7 days

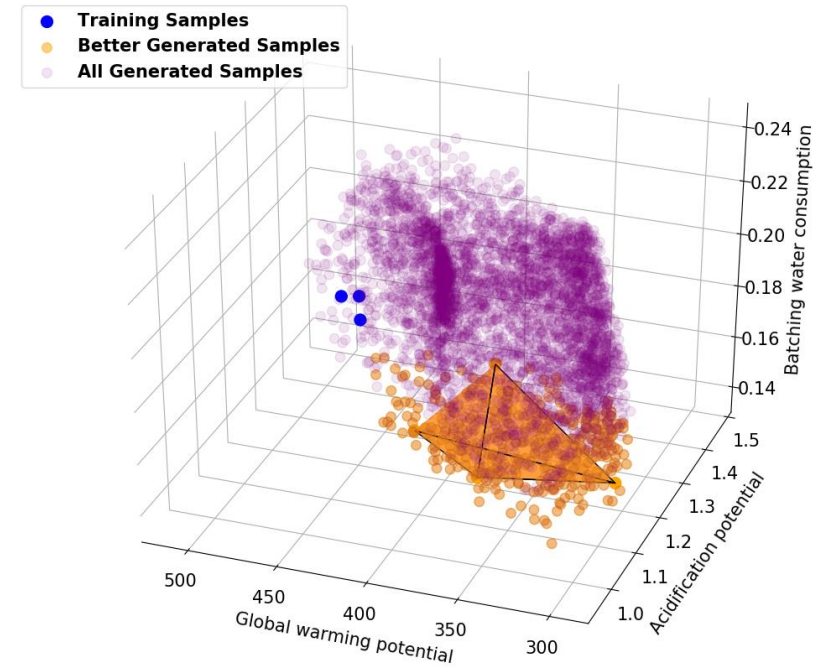


28 days

Environmental impact of training samples and generated samples



Curing time = 7 days
Strength = 30 ± 1 MPa



Curing time = 7 days
Strength = 40 ± 1 MPa

Average environmental impact reduction achieved and newly generated mix designs

Conditional Average Environmental Impact Reduction of Better Samples				
Age	Strength (MPa)	GWP (kg CO ₂ eq.)	AP (kg SO ₂ eq.)	CBW (m ³)
≤3	30±1	0.80	1.83	5.47
	40±1	7.74	1.59	0.26
7	30±1	19.69	3.94	7.58
	40±1	25.45	11.33	5.03
14	20±1	2.20	5.72	10.64
	60±1	42.45	21.09	5.17
28	70±1	21.62	6.66	3.32
	80±1	27.44	8.40	4.15
56	40±1	4.38	2.95	7.04
	50±1	14.38	3.23	3.64
	70±1	30.26	23.75	1.32
	80±1	5.88	1.33	3.46
≥90	80±1	30.58	6.91	4.11

Examples of Generated Mix Design

Strength (MPa)	30±1	40±1
Constituent Material	Amount (kg per m ³)	
Cement	186.4	259.0
Slag	236.7	288.6
Fly Ash	107.1	58.8
Water	142.3	142.5
Superplasticizer	22.3	26.1
Coarse Aggregate	901.4	868.6
Fine Aggregate	717.2	763.0

AI mixes

All units in (kg/m3) unless specified

Cement	Blast Furnace Slag	Fly Ash	Water	Super- plasticizer (SP)	Coarse Aggregate	Fine Aggregate	Age (day)	Target Concrete compressive strength (psi)
131.46	201.21	119.67	180.7	7.46	950.72	780.48	28	4000
128.59	197.46	124.24	184.31	6.61	954.48	787.47	28	3000
134.89	182.74	113.78	179.43	7.32	953.22	785.28	28	3000
132.25	184.37	119.74	181.03	7.33	954.1	786.55	28	3000
129.02	210.6	122.8	184.63	6.84	953.5	780.11	28	3500

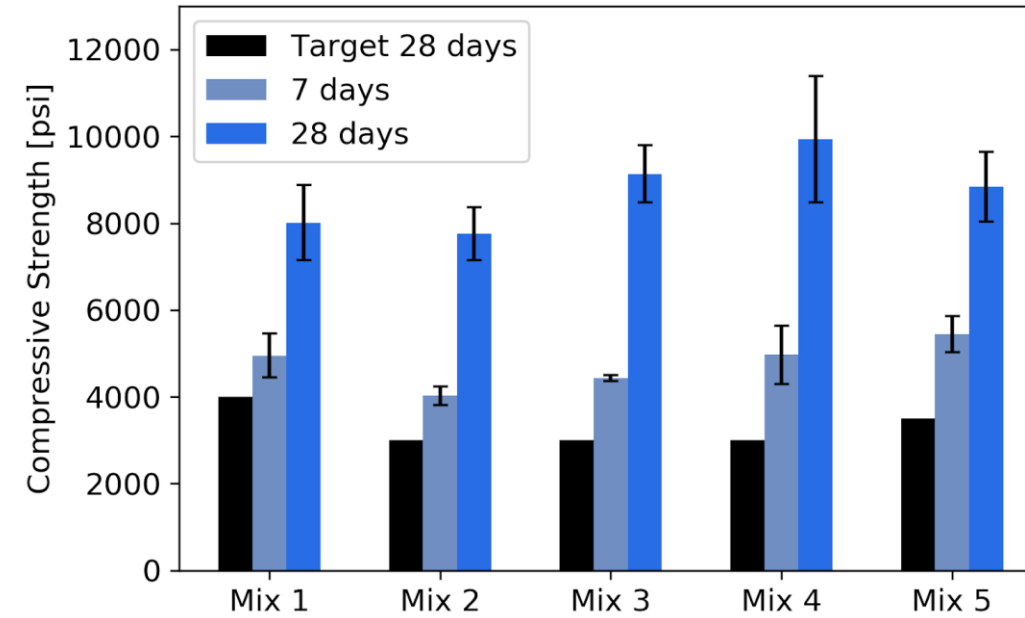
Human adjustment of superplasticizer to improve rheology (and address drift from historical data)



Very good compressive strength in laboratory setting



Compressive strength by mix

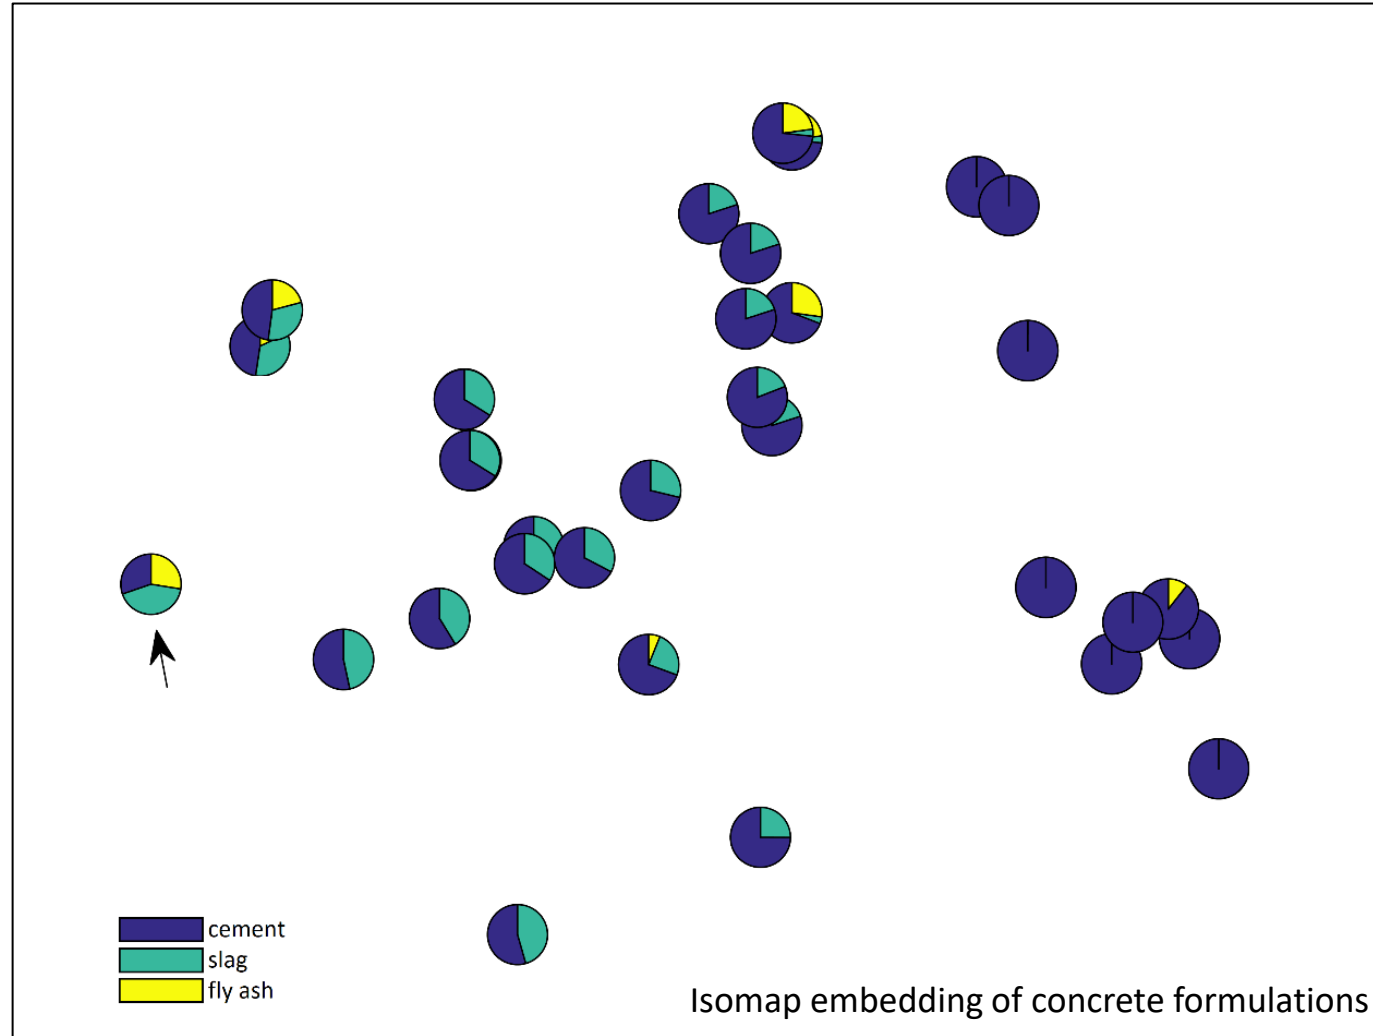


More than 50% reduction in carbon emissions

AI-Based Formulation	1	2	3	4	5
Estimated GWP (kg CO2 eq.)	152.15	152.28	157.29	155.16	152.15
Average of industry standard (similar 28-day compressive strength)	282.36	280.31	318.75	302.45	279.78

Ozinga Formulation	1533SX	1097SX	1109S	1101S	1160S
Estimated GWP (kg CO2 eq.)	207.30	284.92	250.83	249.86	276.33

Extrapolative rather than interpolative



Test Pours at New Data Center

