Chairs, Caricatures and Optical Illusions

Artificial General Intelligence - AGI Digest

Todor Arnaudov

Contents

Todor Arnaudov's theses and analysis expressed during public discussions in the AGI List. Main opponent - Mike Tintner. Other participants:

Piaget Modeler
Ben Goertzel
Sergio Pissanetzky
Matt Mahoney
Jim Bromer
Abram Demski
Charles Hixson

Note: Minor stylistic corrections. May still have some typos, mistakes. Emails are somewhere in the archive of AGI list. Published in this digest by Todor Arndudov in 10/2017.

http://research.twenkid.com

About [7.10.2017] (DD/MM/YYYY)

Todor is the author of probably the world's first University courses in Artificial General Intelligence back in 2010 and 2011 in Plovdiv University, Bulgaria.

Todor is been doing a “rebellious” independent research with most of the work yet in “stealth” mode. Winner of several prizes for ideas and contributions from the fellow indie AGI-er Boris Kazachenko.

31.12.2017

If I’m not mistaken, Geoffrey Hinton agrees with my claim that vision and “invariance” is ultimately 3D-reconstruction, thus “equivariance”, not “invariance” in their words. That’s according to:
– “Does the brain do Inverse graphics”, published on 25.08.2015, a record from a lecture in a “Graduate summer school”, Toronto, 12.7.2012” from
https://www.youtube.com/watch?v=TFIMqt0yT2I

Slides by Kyuhwan Jung, 9/11/2017: ...p.8: “We need equivariance, not invariance”
https://www.slideshare.net/kyuhwanjung/vuno-dl-seminarcapsnetskyuhwanjung20171109

See letter 4 from the “General Algorithms and ...” from 2012 and the appendix: “Colour Optical Illusions are..."

Generalization – Food and Buildings

Thread: General Algorithms or General Programs [was Re: [agi] Unification]

(...)

Mike Tintner <tintner@blueyonder.co.uk> Wed, Jan 11, 2012 at 12:25 PM

Todor:And if you can't generalize the algorithms...

Do you think there *is* such a thing as a "generalized algorithm"? Example? Or do you recognize an unsolved problem?

It's actually a contradiction in terms.

The only generalization of an algorithm which says:

"take THIS step and then take THAT step" [precisely instantiable]

is one which generalizes and says:

"take A step (of a general, nonspecific kind) and then take SOME OTHER step (of a general, nonspecific kind) ..."

That is a truly *general program* for it still gives instructions - but it has ceased to be an algorithm, which always gives specific, precisely instantiable instructions

Only general programs are possible. Algorithms are always specific not truly general.

Todor Arnaudov <twenkid@gmail.com> Wed, Jan 11, 2012 at 7:56 PM
To: agi@listbox.com

Mike,

More general just means covering wider range of possibilities/cases.

Generalized algorithm in the previous context is one that finds analogy not only between word pairs (NLP), but also between visual and sound objects - using the same algorithm.

Mike>take A step (of a general, nonspecific kind) and then take SOME OTHER step (of a
Mike>general, nonspecific kind) ..."

Mike>That is a truly *general program* for it still gives instructions - but it has ceased to
Mike>be an algorithm, which always gives specific, precisely instantiable instructions

Mike >Only general programs are possible. Algorithms are always specific not truly Mike>general.

Todor: Actually algorithms are "general" and abstract, especially if written in higher level language.

And that's your definition of algorithm. I have a term "running/actual algorithm" - that's the sequence of actual instructions/transformations that happened in a CPU or whatever, the real causal chain, not the abstract static code of the program. It's virtual.

In a physical world, all algorithms (everything) are eventually converted into precisely defined executable "instructions" (my term: "machine code of Universe") and physical transformations of their virtual "actual algorithm".

Mike>An algorithm is *totally specific* It has to tell the machine for example which space of possible hand configurations to explore in picking up a block.

Todor: The same goes for humans... You have to give *TOTALLY SPECIFIC* trajectory of your hand (...), it's general only in your head and you don't realize the details of the lower level abstractions - this is the passing from higher level of abstraction to lower, down to the physical level which is the most specific.

-- Todor

Piaget Modeler <piagetmodeler@hotmail.com> Wed, Jan 11, 2012 at 9:59 PM

What if the grasp plan was differentiated into diverse subplans which were each ultimately executed by various algorithms? So we have the general plan to grasp, and then depending upon the context, the robot would select the appropriate algoritm for the kind of object being grasped: a rope, a cactus, a gun, a bottle, etc.

Would that work?

Mike Tintner Fri, Jan 13, 2012 at 1:08 PM (16)

In my modest view of this discussion, case based reasoning provides an answer to the open-endedness of the "world of options". We just need to have a set of cases that is sufficient to reasonably handle the problems encountered. The case base can incrementally expand to include new, unanticipated scenarios. But I don't think the human or artificial mental case base needs to be as infinite or as wide ranging as the number of options in the world in which a person or an AGI would function.

PM, (Matt?!)

This is only a first tentative step to answering your point – but I think it's basically absurd (although clearly seems reasonable to you & others) – and **one ought to be able to do the maths here *** (wh. is why I'm appealing to Matt, Ben et al).

Basically, what you're saying comes down to :

there is or can be a set of variable recipes [/sets/programs] for all foods, all dishes, all menus and all forms of preparing food

[or there can be a set of variable blueprints[/sets/programs] for constructing all buildings]

Well, first off, there can't be - because the world and all its fields and domains, incl. food and building, are continuously evolving, and you can't tabulate the new.,

Secondly, you can't even define any general field or domain, like "food" or "building". That's a huge and ultimately impossible creative job in itself.

But even if you could set arbitrary limits to a domain, the maths would be unfeasible - your recipes/programs would quickly become impossible.

And that's where the maths - and plea for help - comes in. This is after all the problem of complexity that fascinates so many here.

Let's say we define a set of a million foods.

How large would the set of rules and program have to be, in order to define all the possible combinations of those foods?

(I don't think you could even define the endless ways of combining/ cooking/mixing/slicing and dicing etc foods together, let alone the possible styles of menu.)

But again, set arbitrary limits to the ways of combining foods.

How large a program would you need?

And contrast this with chess, where all the combinations of just a few pieces on a few squares, is still beyond the compass of present programs/computers.

Come on, mathematicos.

Piaget Modeler <piagetmodeler@hotmail.com> Fri, Jan 13, 2012 at 8:46 PM

Mike Tinter:

This is only a first tentative step to answering your point - but I think it's basically absurd (although clearly seems reasonable to you & others) -
and *one ought to be able to do the maths here* (wh. is why I'm appealing to Matt, Ben et al).
...
Well, first off, there can't be - because the world and all its fields and domains, incl. food and building, are continuously evolving, and you can't tabulate the new.,
Secondly, you can't even define any general field or domain, like "food" or "building". That's a huge and ultimately impossible creative job in itself.
But even if you could set arbitrary limits to a domain, the maths would be unfeasible - your recipes/programs would quickly become impossible.
And that's where the maths - and plea for help - comes in. This is after all the problem of complexity that fascinates so many here.
...
Come on, mathematicos.

Confucius says that man who says it can be done, and the man who says it cannot be done, are both usually right.

~PM

Todor Arnaudov <twenkid@gmail.com> Fri, Jan 13, 2012 at 9:34 PM

Guys,
I've not read all the latest bunch of comments, regarding that last one cited by PM is pretty easy to me:

-- Noise is potentially infinite, also - meaningless.
-- Meaningful regularities in real world in limited space-time and discoverable by limited resources are finite, though.
-- In your specific example of "food" - it's a word, words are abstract and ambiguous, meaning they have different more specific definitions in different contexts - not big deal, you have to provide context, and in real world it's always there.

Let me give some sample actual (generalized) meanings of "food", based on very easily detectable and possible to generalize by sensory experiences. Of course these are not complete definitions, it's the principle:

A) Something that when put in the mouth, chewed, swallowed etc. after some amount of time causes *increase of the glucose level* or causes it to stabilize. (That means, if nothing which is food is put in mouth, glucose level decreases, or increase in cortisol is detected, or other metabolic changes happen.
In more "scientific meanings", food is a matter that turns into ATP or initiate metabolic processes which turn already stored biological matter to break down to ATP, but this is a very abstract.

B) A plant (an entity which is known to grow out of soil - has green tissues, or it (the fruit or vegetable) are connected with entity which is green, flexible,... -- these details can be taken by incremental learning and matching a fruit on the table with fruits in pictures or reality, drawn connected with the trees.)

C) An animal (grows, has soft tissues, moves; is similar to another one which has these features; is similar means that some of their features match in particular level of abstraction, while they don't match if compared to other concepts ...)

D) Particular tastes, associated with eating some of the above.

For further specifications and groups: "dairy", "sweets", "pastry", "chocolates" -- all the data to group them is there without supervision, such as:

-- Color, particular consistency, smell, taste.
-- Known to be made out of something (visual appearance, or knowledge of visual appearance before cooking etc.
-- Etc.

As of "building" -- what's so special?

You're looking at the non-essential features, they can be endless amount, that's why they seem "infinite" to you.

Building is generally something that's built - whatever the details or the materials. (Again it's a matter of context an *choice* what to be assumed as "legal" building, such as is a "paper house a house", but that's not essential)

"Built" means - it didn't exist before the event of "building".

The "building" as an event is *causing transformations* which are altering environment so that the spatial and bordering (limiting) characteristics on the location where this process is happening are changing.

Such transformations are:

-- Moving (translate and rotate) of an entity -- meaning - change of its physical coordinates, in sensory terms -- changes in the recovered coordinates in 3D-space.
-- Morphing an entity - applying controlled directed forces which alter the geometrical features of an entity to a desirable shape, such as modeling bricks out of clay.
-- Gluing/fitting together/constructing -- meaning altering the spatial and physical relations between entities so that they are now physically (and visually) adjacent (touching) and the force which is expected to be needed to unglue/break the bond is bigger than it was before the process of gluing.

Etc....

-- The Building itself is:

Something (a discrete area in space, which can be recognized as separated from the rest) that has
-- a COVER (something which would keep the area below dry if it's raining above it)
-- a WALL (something which would stop something that moves towards it, such as wind).

All the rest are details which can be there, or cannot be there. Even the two items above are not obligatory to happen in the same time, e.g. -- a building *in progress* which is lacking a roof.

-- Todor "Tosh" Arnaudov

Mike Tintner <tintner@blueyonder.co.uk> Fri, Jan 13, 2012 at 9:42 PM

Todor: As of "building" -- what's so special?

Please explain how you are going to define the space and potential forms of possible modern architectural buildings, using recent third millennial buildings as a guide:

http://blog.360dgrs.nl/wp-content/uploads/vilnius3.jpg
http://techandsoc.com/files/2011/02/Guggenheim-Bilbao-2.jpg
http://www.cctv.com/newSiteProgram/images/project/1.jpg
http://images.businessweek.com/ss/07/06/0628_chad/image/slide-2.jpg
http://4.bp.blogspot.com/-_XAO7CTQhYA/Tie8kyxPa0I/AAAAAAAAB_U/qvxuhkwM4jo/s1600/dynamic-modern-architecture-design.jpg
http://greendesign1.com/wp-content/uploads/2011/04/Modern-Architecture.jpg
http://www.besthousedesign.com/wp-content/uploads/2011/03/Modern-Architecture-by-Emmanuelle-Moureaux-Architecture-+-Design-1-588x570.jpg

{And since you're a fan , you might try explaining how you would apply Boris' pattern discovery principles here to help you).

Mike Tintner <tintner@blueyonder.co.uk> Fri, Jan 13, 2012 at 10:16 PM

I think to answer myself that what my examples show is that buildings can take an infinite diversity of forms - albeit bounded by provisional constraints. You couldn't have a floating house for example (at least not for a while).

And if you think about it, food can take an infinite diversity of forms too.

And that's before we even start thinking about possible materials or ingredients, chemicals, and nano/atomic/molecular etc rearrangements - let alone all the possible modes of construction and cooking.

Confucius should have followed his own advice and thought twice before speaking.

P.S. It would be good to provide some formal demonstration of the infinite creative possibilities of virtually any medium - because there are so many misguided people out there thinking that formulae/algorithms/recipes can comprehensively embrace everything under the sun

My guess is it might lie in: how many different lines of any shape and length can be drawn on an A4 page?. I would guess there's an infinity but it would be good to prove it one way or the other. Could that be a Clay candidate, Matt? Perhaps the recent demonstration that a given rope can be of infinite length once you consider its fractal nature/edges might be relevant here.AGI | Archives | Modify Your Subscription

---

Ben Goertzel <ben@goertzel.org> Fri, Jan 13, 2012 at 10:19 PM

Can you define what you mean by "infinite diversity" ?

Or do you consider that concept, in itself, undefinable?
[Quoted text hidden]
--
Ben Goertzel, PhD
http://goertzel.org

"My humanity is a constant self-overcoming" -- Friedrich Nietzsche
AGI | Archives | Modify Your Subscription

Mike Tintner <tintner@blueyonder.co.uk> Sat, Jan 14, 2012 at 12:20 AM
Reply-To: agi@listbox.com
To: AGI <agi@listbox.com>
Ben:Can you define what you mean by "infinite diversity" ?

In this case, I guess it's any form - including any line - that is not a variant of the same pattern as another, and therefore unique.

So you could obviously draw a range of simple wave lines, which would simply be magnifications of each other They are not diverse - and count as just one kind of form. Vertical lines of different lengths are similarly not diverse - just one form. But the slightest irregularity - the slightest nick in a straight line would make it an individual, diverse.form.

By all means sharpen up the definition. The objective here is among other things to explore the limitations if any of design space. What are the limitations if any of designing tools and buildings (considered just as 2d-forms on that page)?

I've just BTW been getting into the latest freeform architecture (v. much spiritually in tune with my ideas re AGI) and I'm not usually that enthused, but I find it mindblowing beautiful stuff. (These are all individual, diverse forms).

http://www.google.com/search?hl=en&safe=off&biw=1446&bih=854&gbv=2&tbm=isch&q=zaha+hadid+buildings&revid=1634987981&sa=X&ei=uKsQT8yUCYnu8APY7vTtAw&ved=0CEAQ1QIoAQ

http://www.evolute.at/consulting/references/yas-island-racetrack-hotel.html

http://www.google.com/search?tbm=isch&hl=en&source=hp&biw=1446&bih=854&q=freeform+architecture&gbv=2&oq=freeform+architecture&aq=f&aqi=g1g-mS1&aql=&gs_sm=e&gs_upl=2965l7293l0l7662l21l19l0l1l1l0l234l2210l3.11.1l15l0http://www.evolute.at/consulting/references/yas-island-racetrack-hotel.html

---

Ben Goertzel <ben@goertzel.org> **Sat, Jan 14, 2012 at 12:38 AM**
Reply-To: agi@listbox.com
To: AGI <agi@listbox.com>

In that case, I don't see why you think existing AGI algorithms are in principle in capable of handling "infinite diversity"

For instance, an automated program learning algorithm like MOSES or Genetic Programming can -- in principle -- recognize or create an infinite diversity of forms, in your sense... so can a recurrent neural net (in all cases: assuming adequate computing resources)

ben[Quoted text hidden]
--
Ben Goertzel, PhD
http://goertzel.org

"My humanity is a constant self-overcoming" -- Friedrich Nietzsche
AGI | Archives | Modify Your Subscription

Todor Arnaudov <twenkid@gmail.com> **Sat, Jan 14, 2012 at 2:07 AM**
To: agi@listbox.com

MIke,
GI doesn't need to define the complete *"space and potential forms of possible modern architectural buildings, using recent third millennial buildings as a guide"* or neither medieval ones.

The Space of Buildings is: a Wall (and/or) a Cover (Roof)

In order to answer the question is something a building, GI only needs the patterns required to recognize that a "building" belongs to a class of buildings (and as I noted usually such abstract

concepts can be classified in many classes depending on the selected features, if selection is arbitrary).

The process above requires particular **context (experience, records)** having samples called "building" where the GI induces what **all these samples have *in common*, not how they're different.**

I see you don't get the "**generalization**" (…)
You're searching for "the most different things" - they don't make sense.

Intelligence is in finding the common between things and simplifying them (...) - compressing. Combinatorial explosion is inevitable if raw sensory input is not reduced and compressed.

In this case and those pictures:

All those things ***touch*** the ground in the 3D-reality reconstructed from the images and are going higher than the ground, and they are not rocks or earth (because of their visual appearance), the rocks, earth, trees etc. are not "buildings", because they're not constructed (built). "Construction" of "buildings" in that context is something attributed to "agents", geological or different-than-human processed are not considered as such.

In general, all the objects on the pictures seem to have "roofs"/"covers" and/or walls - like in my previous definition. (Windows are also "walls" in that definition, there are buildings with no windows and with no doors, that is not essential.) This is the definition of the space, simpler than you thought... :)

Another feature is their size - they appear big enough to hold people inside. (But it's not always relevant, there are model buildings)

They can be floating, not touching the ground - OK, a more abstract meaning of "ground" is "that thing below the other things which are touching it and direct in the opposite direction".

Bottom line: all those don't matter:
-- ...the materials used in the building, the very specific shape, the material of the ground, the sizes, the colors, the age, the lightning of the picture, the resolution, the angle of the camera etc.

The features above are "specifics", namely because there are many different samples where they hugely differ, they apparently cannot be generalized to a common class, because they are different.

If the focus is more specific, the buildings can be spread in various different classes, though:
-- With windows - mirror from the external side or not
-- Using steel
-- Oval shaped
-- "Angularly"-shaped
-- Skyscrapers
-- Towers
-- Twin Towers
-- Urban
-- Surrounded by natural landspace
-- ... Whatever

Those are "**different classes of building**".

Of course, for the very specific shapes there can be defined 3D models using various resolutions, blueprints covering many enormous sheets, and detailed definitions of every bit of the building.

All those are **not just "buildings"** these are ***specific cases of designs of buildings*.**

They can have various resolutions - so there can be "low" or "higher resolution definitions of designs of "buildings"".

Abstractions for Dummies

"Building" is an abstract word, "abstract" means namely lacking the specific details and bearing only the general, essential.

As of Boris' principles, I share them, e.g.:

-- Generalization (explained, and I told you what's general even in the previous email)
-- Compression (explained)
-- Prediction (that's more for the process of building the hierarchy and feedback, and the process of recognition of the patterns as belonging to a particular class more than to another)

Features/details/specifics which are too different/unpredictable are those which are too specific given resolution of perception and prediction - can't be predicted correctly by the model with sufficient precision that the model aims at the particular level of generalization -
are pruned from the description vector, only the essential ones - the general - are kept, otherwise the system will go out of resources.

Another example -- it won't be very handy if we call each building differently and tell all minute details all the time, our mouths are too slow. Like:

"Hi honey, do you like this 135-m-tall-the-basis-is-10.456mx9-457m-the-first-window-has-the-following-shape-ellipse-center-..there's a table on the 2-nd floor, it's coordinates are ... it's size is... colour texture -... the light pattern ... .-radius-the-next-window-... (...) -- the last floor-it.s.."

...and every time the other person would have to filter all that **JUNK** data until she discovers that this is about something that's of her interest.

One could also speak out loud every single pixel of a photograph she has taken:

- Hi honey, do you like us to go this hotel for the weekend, I got three options, they look like this:":

Speaking out loud:

640x480 RGB (to be modest), Bottom-down, horizontal-to-vertical: 0xff00f3, 0xff5654, ... (...) 0x5644223 ...

1024x768 RGB, ...

2048x1536 ... 0x4454534 0x3462384

- Which one of those do you like best as our home?...

Obviously that's not GI, but a **literal data transfer.**

Of course, in order to communicate specifics which are hard to generalize, human GI use low level sensory records or models which are closer to the real source, such as images (3D to 2D projection), sets of images, videos, blueprints -- as you did by giving pictures as samples.

The reason is precisely because those details are too specific to be effectively generalized, it's more effective to transmit them on demand on the lowest levels only.

Bottom Line

In fact, the bottom line and the essence of your "complains" is that natural language is not good enough to communicate specifics which require images...

Well, GOOD MORNING!

Natural language - if derived from sensory experience and grounded - is an excellent example of GI.

General - from *generalization*. The patterns it uses are *abstract*, meaning - the irrelevant details are lost, they're left in the *lower level representations*, which are accessible by the higher levels, however the patterns in abstract terms and specific terms are not completely comparable - generalization is lossy compression.

-- Todor "Twenkid" Arnaudov

(...)

Charles Hixson <charleshixsn@earthlink.net> Sun, Jan 15, 2012 at 9:40 AM

No. That's not correct. It's just that the generalization of an algorithm is likely to be incorrect.

E.g., on could generalize:
To put the fire out take a pail of water and throw it on the fire.
to:
To put the fire out, take a pail of liquid and throw it on the fire.

This is clearly a generalization, but if the liquid happens to be gasoline, it doesn't work very well.

Todor Arnaudov <twenkid@gmail.com> Sun, Jan 15, 2012 at 2:46 PM

Hi Charles,

Charles>To put the fire out take a pail of water and throw it on the fire.
Charles> to:
Charles> To put the fire out, take a pail of liquid and throw it on the fire.

Well, if generalization is done randomly, there's various % chance for correctness, depending of the number of (possibly yet unknown) related classes of lower generality involved.
In this example it's about: is a liquid ***inflammable*** if it's thrown into fire.

If you have thrown just water, and you haven't ever seen (or tested) other inflammable liquids,

you apparently can make that mistake, and this won't be "wrong", you make a guess and may be correct or not, the next time you will know.

If one has enough experience, she may make the right generalization from experience though:

Like

-- Movies with vehicle fuel burning and exploding (and has experience that the fuel is liquid)
-- "Cocktail Molotov" sequences
-- Alcohol burning (for cooking)
-- Even just a pocket *lighter* (if the fuel tank is not transparent - shaking it may suggest it has liquid, also holding bottles with gasoline for lighters.)
-- A more advanced generalization (not critical) - a candle burning, the liquid paraffin under the flame.

Experience shows that *some* liquids can be used to put the fire out, and others - to set fire, so the feature "liquidity" is not enough to decide this, meaning:

--> This degree of generalization is too lossy for the purpose, more specifics are required.

-- Todor "Tosh" Arnaudov

Todor, 16-7-2013

If **it**'s too "free-form"-y and **it** allows too much variety, then the matter (it) gets too general – in the bad sense of general – vague. Then almost everything would pass as "it".

General Algorithms or General Programs [was Re: [agi] Unification]

file://localhost/C:/Users/tosh/Downloads/to_1-9-2012/What's-a-chair-AGI-list-Gmail%20-%20[agi]%20Vision%20and%20causality...%20a%20spinoff%20from%20Further%20uniform-izing%20DeSTIN%20(to%20get%20rotation%20and%20scale%20invariance%20for%20free).htm

file:///C:/Users/tosh/Downloads/to_1-9-2012/What-is-a-building-generalization-Gmail%20-%20General%20Algorithms%20or%20General%20Programs%20[was%20Re%20%20[agi]%20Unification].htm

Generalization – confusion ….

General Algorithms or General Programs [was Re: [agi] Unification]

Discussion in the AGI List: … April 2012 with (in ~ order of involvement and then chronologically):

Todor Arnaudov
Mike Tintner
Piaget Modeler
Ben Goertzel
Sergio Pissanetzky
Matt Mahoney
Jim Bromer
Abram Demski
Charles Hixson

---

(1)
Sergio Pissanetzky <Sergio@scicontrols.com> Thu, Apr 26, 2012 at 6:38 PM
Reply-To: agi@listbox.com

Ben,

I have renamed my thread so it does not interfere with DeSTIN or OpenCog.

I was trying to understand causation in vision by learning from the human vision system.

Evolution has had time to fine-tune it, so it is still worth learning from it.

And I am concerned that rotation and scale invariance are only two of an enormously large number of invariants needed for actual image recognition. The theory of invariants is in its infancy.

Sergio

---

(2)
From: Ben Goertzel [mailto:ben@goertzel.org] Sent: Wednesday, April 25, 2012 6:25 PM
To: AGI
Subject: Re: [agi] Further uniform-izing DeSTIN (to get rotation and scale invariance for free)

Sorry if I wasn't clear -- I'm not trying to make DeSTIN into a better model of human visual processing, rather into a better vision system, and more compatible with OpenCog...

ben g

...

---

(3)
Ben Goertzel <ben@goertzel.org> Thu, Apr 26, 2012 at 7:37 PM
Reply-To: agi@listbox.com
To: AGI <agi@listbox.com>

Rotation, scale, translation and shear are the only linear transformations you need to generate all 2D linear transformations, so they seem to have a privileged position

My suggestion is to handle these basic linear transformations within the low-level perceptual machinery, and to handle more complex invariances via learning...

ben
[Quoted text hidden]
--
Ben Goertzel, PhD
http://goertzel.org

---

(4)
Todor Arnaudov <twenkid@gmail.com> Fri, Apr 27, 2012 at 1:12 AM
To: agi@listbox.com

I don't know if anyone on this discussion realized, that "Invariance" in vision is actually just a

- 3D-reconstruction of the scene, including light source and the objects

- Also colours/shades and the textures (local/smaller higher resolution models) are available (for discrimination based on this, may be quicker/needed for objects which are otherwise geometrically matched)

[+ 16-7-2013 - conceptual "scene analysis", "object recognition" involves some relatively arbitrary, or just flexible, selection criteria for the level of generalization for the usage of words to name the "items" in the scene. To Do: devise experiments with ambiguous objects/scenes, sequences. … see "top-down", … emails 9, 14, 15]

If the normalized 3D-models (preferably to absolute dimensions), lights and recovered original textures/color (taking into account light and reflexion) are available, everything can be compared perfectly and doesn't require anything special, and no "probabilities" or something. The textures and light most of the time don't even alter the essential information - the 3D-geometric structure.

"2D" is just a crippled 3D

"Invariants" in human fully functional vision are just those 3D-models (or their components, "voxels:) built in a normalized space,the easiest approach for quick comparison is voxels, it might be something mixed with triangles, of course textures and colours also participate.

Every 3D-model has a normalized position per its basis, and also some characteristic division of

major planes and position between the major planes, and there are "intuitive" ways to set the basis --> **gravity/the ground plane foundations, which is generalized to "bottom", i.e.:**

-- The "bottom" of an object, which faces the ground, is the part of the image of the object which projects on the "bottom" of the scanlines of the retina, because that's inferred for the first objects, which always have stable touch with the "ground".

When generalizing or specializing, the resolution of the 3D-models to be compared is changed (see the thread where I gave example of how the concept of a "building" is produced), at particular stage every two 3D-models match, eventually converge to a cube, or a plane.

IMO in fact brain is not very good in further mental rotation of those models, yeah we know those IQ tests, but humans do it very slowly and the tests consist of very few crossing planes, because it gets too complex.

Con: "How can you say that it's "just" 3D-reconstruction? That's so compex!"

- Well, one may think so only if she was not familiar with the triangulation (photogrammetry dates back to 19-th century) and/or the spectacular work of Mark Polleyfeys.

"How do you recognize that this is your chair, if it's upside down and you haven't seen it before"

Like the mistakes about generalization - a "chair" is a generalized concept, it's not a pixel-by-pixel image, rough 3D-models are compared for finding a match. And matching is a biggest number of high degree of match of the size relations of the boxes, planes color (after light correction) + texture, to the match to those of the chair from the previous day, than to those boxes, planes etc. of "chairs" found elsewhere, and of any other "objects".

A "**chair**" **[a stool]** generally is just:

-- A **plane** which is **perpendicular** to the "**ground**" **direction vector**, which is a vector which is parallel to "gravity" - that is the vector where objects go when let without a support;
-"**support"** is a vector consisting of "solid" connection (of forces, impacts) to the "ground" which when existing prevents objects from getting closer to the "ground" (falling);
- **the "ground"** is a plane where objects **stop their motion** (changes of coordinates between subsequent samples) if left without support or impacting by other moving "things", etc.

Most chairs can be reduced to a few solids and still be recognizable.

AGI is way simpler than it seems.

-- Todor "Tosh" Arnaudov
http://research.twenkid.com
http://artificial-mind.blogspot.com

Matt Mahoney <mattmahoneyfl@gmail.com> Fri, Apr 27, 2012 at 5:15 AM (5)
On Thu, Apr 26, 2012 at 6:12 PM, Todor Arnaudov <twenkid@gmail.com> wrote:
> A "chair" generally is just:
>
> -- A plane which is perpendicular to the "ground" direction vector, which
> is a vector which is parallel to "gravity" - that is the vector where
> objects go when let without a support; "support" is a vector consisting of
> "solid" connection (of forces, impacts) to the "ground" which when existing
> prevents objects from getting closer to the "ground" ; the "ground" is a

> plane where objects stop their motion (changes of coordinates between
> subsequent samples) if left without support or impacting by other moving
> "things", etc.
>
> Most chairs can be reduced to a few solids and still be recognizable.
>
> AGI is way simpler than it seems.

No it's not. http://bit.ly/IIPMuz

-- Matt Mahoney

[Pictures of charis from google search, e.g.:]

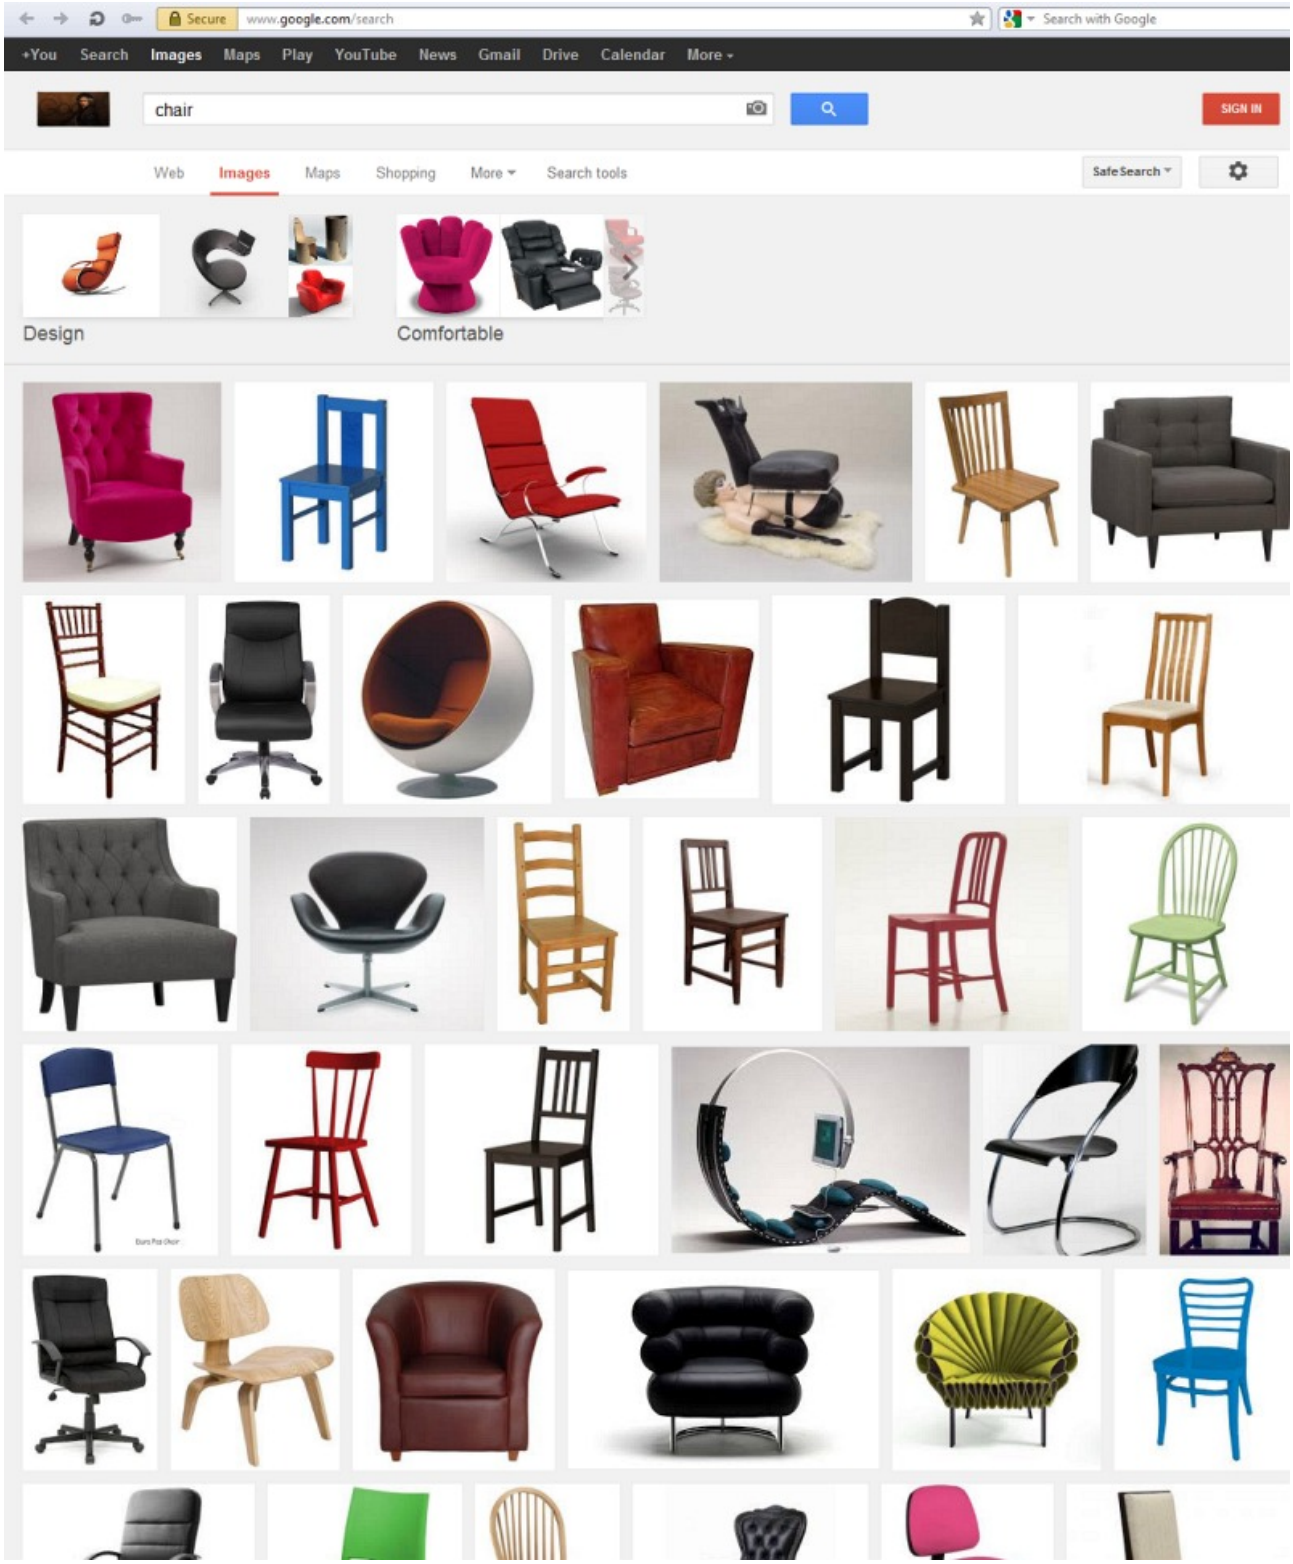

Todor Arnaudov <twenkid@gmail.com> **Fri, Apr 27, 2012 at 10:46 AM** (6)
To: agi@listbox.com

Matt,

Too bad for you that you don't understand it, that's the same example like the "buildings" of Mike.

All those pictures comply with the above definition (...), and it's the **essence (generalization) of**

all samples, after 3D-reconstruction and light-source recovery are done.

3D-models are invariant per-se, to their basis of the coordinate system, then matrices of transformations are applied.

In modern GPUs there's a feature called "instancing" which is that a model is populated at different coordinates, but the 3D-description is stored and read only once, just the difference vectors are used.

...

A chair must have just this in order to be a "chair":
[I've meant the general of "chair" and "stool", more precisely – in Bulgarian both are "chairs"]

[CHAIR] -- A plane perpendicular to the "ground" and parallel to the gravity vector (in its canonical/normalized form). That plane is supported by something (a vector (of force, might be flying), mesh; or a curve that may have a normal which is perpendicular to the ground) etc.

All the rest of the features: the back, elbow-rest, colour, size, number of legs, *lightning* of the scene, rotation, scale, position on the viewport - these are **decorations** and **specifics**, they don't matter at the level of abstraction for recognizing just is it a "chair" - just one appropriate perpendicular plane + support have to be found.

The "**decorations**" are **for recognizing different types of chairs**:
- three-legged
- four-legged
- with back
- without back
- red, green, [etc. color]
- with elbow-rest
- big or small
- turned upside-down [or in "cannonical position"]* etc.

The principle is the same, just there are more planes/solids to compare, and also colours, textures (small pictures which are following the planes/curvatures of the geometry of the object).

Cheers,

-- Todor "Tosh" Arnaudov
http://research.twenkid.com
http://artificial-mind.blogspot.com

[Quoted text hidden]

Mike Tintner <tintner@blueyonder.co.uk> Fri, Apr 27, 2012 at 10:26 AM (7)

Matt: >> AGI is way simpler than it seems.

No it's not. http://bit.ly/IIPMuz

And those, if you remember my chair post/illustration, are the **easy** ones!!

Mike Tintner <tintner@blueyonder.co.uk> Fri, Apr 27, 2012 at 11:27 AM (8)

Todor: Matt,,Too bad for you that you don't understand it,

So... explain on Matt's page how chairs on

Page 1

Line 1 ... Chair 4 across (arm rests, solid back) becomes/is transformed into Chair 5 across (no arms, multilinear back) .. .. and then into...

Line 4.. Chair 2 across {"coiled" chair - no arms) is transformed into Chair 7 across (Mirra chair - mobile wheels, armrests) .. and then into

Page 2

Line 1.. Chair 2 across (Ron Arad chair -undulating form)

Line 5... Chair 6 across (solarsystemchair - consisting largely of balls)

Explain how the concept of "chair" embraces all these radically diverse geometric forms, and radically diverse selections of parts - using your 3d principles.or any geometry you like.

The "Principles of transformation" involved in the class of chairs cannot be embraced by any existing geometry. "Freeform geometry" points the way to understanding - but at the moment the freeform transformations are generated by the geometer not the geometry - as so often in AGI matters.

Todor Arnaudov <twenkid@gmail.com> Fri, Apr 27, 2012 at 12:24 PM (9)
To: agi@listbox.com

Mike,

The concept doesn't emerge from all those chairs - those chairs are compared and reduced to the concept.

The word "chair" is probably thousands of year old, and in fact initially there were just one (for an individual child) or a few types: of chairs [places where one could sit, stools]:
-- a stone
-- a stump/stem
-- a step-like piece of earth

Little children can grasp the concept of "chair", because they see the simplest in it (as in my definitions), they can't know all kinds of "chairs" and don't need that. **The function of the chair is to sit, which requires just a plane, perpendicular to the gravity, having a support.**

The other samples are reduced to the initial models, created from the first samples.
(They are just reduced, but at maximum reduction they match most with "chairs", and not with "wheels" for example.)

The basic objects are just the basic partitions of space into simplest kinds of 3D-models, having a few number of planes, solids, perpendicular and parallel planes, curves etc.

Then the objects (3D-models) encountered later are compared to those basic partitions.

That's how the older 3D-games and 2D-games still worked, even for Atari 2600 -- **a "human" is recognizable either with a very few pixels, solids, vertices, boxes, curves, or with millions of them in 3D with correct light** - because the millions can be reduced to a few by simple averaging, "triangles reduction".

An example, from 13 million to 400 faces:
http://www.atangeo.com/examples/horse

A recognizable human that can "walk" can be represented by <20 boxes.

>Page 1
>
>Line 1 ... Chair 4 across (arm rests, solid back) becomes/is transformed into Chair 5 across
>(no arms, multilinear back) .. .. and then into...
> Line 4.. Chair 2 across {"coiled" chair - no arms) is transformed into Chair 7 across (Mirra
>chair - mobile wheels, armrests) .. and then into
>Page 2 - Line 1.. Chair 2 across (Ron Arad chair -undulating form)
> Line 5... Chair 6 across (solarsystemchair - consisting largely of balls)
>Explain how the concept of "chair" embraces all these radically diverse geometric forms, and
>radically diverse selections of parts - using your 3d principles.or any geometry you like.

I already did, let's repeat for the students who are not fine yet:

1. Reconstruct 3D-model of the object, the direction of gravity, the plane of "ground".
2. Reduce the resolution of the model (number of vertices, planes, voxels, boxes, solids), **until the point when those diverse models match.**
3. The **reduced resolution representation is the abstract concept of all those samples.**

"Balls" -- if the representation is in voxels with low-enough resolution, there are no balls, just voxels - space which is occupied, vs. space which is not occupied, it all turns into a number of filled/empty boxes.

>The "Principles of transformation" involved in the class of chairs cannot be embraced by any
>existing geometry. "Freeform geometry" points the way to understanding - but at the moment
>the freeform transformations are generated by the geometer not the geometry - as so often in
>AGI matters.

Those "principles of transformation" are confused, those transformation are as you say of the materials and parts to produce chairs. The recognition of the common between these different "transformation" is possible because of the **reduction to small number of countable features which do match**, in the case of "chair" **it's also a top-down approach:**

-- Mind *searches* for a "chair". It's in fact just a picture or a 3D-model, the purpose may be different etc., the mismatching features can be infinite and random. Mind deliberately reduces the number of features to match, in order to put the samples into a common class.

-- Todor "Tosh" Arnaudov
http://research.twenkid.com
http://artificial-mind.blogspot.com

Mike Tintner <tintner@blueyonder.co.uk> Fri, Apr 27, 2012 at 12:54 PM (10)

Maybe Matt can explain to you that what you're saying doesn't compute, because I doubt I can.

If you want to make this v. simple, concentrate on just one area of a chair - explain how 4 legs [classic wooden chair] become the single folded plane of the Ron Arad chair quoted - their "3-d models" no matter how reduced, how voxelised, will never match.

You do not manifest any understanding that visual object recognition is an unsolved problem - which happens to be a sociocultural fact. The example I have just given you crystallises why it is.

Classes of real objects embrace radically diverse forms, structures and parts - and quantities of parts . For example, I can design a chair composed entirely of snakes - literally, plastic snakes. I can also design a chair composed entirely of triangles and balls. Triangles and balls are not snakes - there is no geometric operation that can recognize one as the other. But you are in the position of claiming there is - and this is an "easy" problem.

If it's easy, you should have no problem explaining how 4 legs become a single folded plane - concentrate on that. But no waffle - directly explain their equivalence.

P.S. One interesting thing that is emerging for me from this is that what appears to be a question of comparing separate *objects* is vastly more complex than that. We are all thinking about chairs, cars, houses and the like as if they were just isolated forms, 2-d,3-d or otherwise. In fact, the recognition of those objects is **inseparable from their relationship with human beings** - with what human beings do with them. The business of recognizing involves comparing and aligning objects like, say chairs, with human forms. We are not consciously aware of that part of recognition most of the time, but it is there. There is a lot more here than meets the eye (and what meets the eye is hard enough as it is).

Todor Arnaudov <twenkid@gmail.com> Fri, Apr 27, 2012 at 4:27 PM (11)

Mike,

I told you the human-related reasons for a chair to be a plane, and how this emerges from experience - **such as something to sit on and not to fall, the simplest one is what I defined.**

>You do not manifest any understanding that visual object recognition is an unsolved problem -
>which happens to be a sociocultural fact. The example I have just given you crystallises why it is.

That's one of the shames of CS, your example demonstrates your **incapability to understand the process of generalization, like earlier.**

[##RADICAL BULLSHIT]

>Classes of real objects embrace radically diverse forms, structures and parts - and quantities
>of parts . For example, I can design a chair composed entirely of snakes - literally, plastic
>snakes. I can also design a chair composed entirely of triangles and balls. Triangles and balls
>are not snakes - there is no geometric operation that can recognize one as the other. But you
>are in the position of claiming there is - and this is an "easy" problem.

It is an easy problem, I'll give the "formal" prove once I start to really work on my AGI prototypes. **You mess up abstract concepts with specific concepts** [even specific entities, not concepts; single, specific instances, raw records, raw sensory data], **it's not your fault only, I see that most people do, also in AGI.**

[##*Note from 20/7/2013*:
People do generalize and recognize generalizations, but they **can't reason** about the generalization process.
- **They see only the results.** "Is it from the class or not."
- They **mix specifics with the general.**
- Can't intentionally operate with **different resolutions –** get an input that's one resolution, for example let's take a chair, made "of snakes", then **"melt it down"**, so that the snakes can't be recognized anymore, and then "crystallize" it again. VoilA!

***** Имаше някаква научна статия за мозъка ли, за какво беше, дето беше объркано обобщаването, същата грешка – и статията на Хофстадтер за шрифтовете – да намеря!!! ... Някъде из любими ...?**
]

>If you want to make this v. simple, concentrate on just one area of a chair - explain how 4 legs
>[classic wooden chair] become the single folded plane of the Ron Arad chair quoted - their "3-
>d models" no matter how reduced, how voxelised, will never match.

Four legs, 100 legs – the count doesn't matter - legs are just **an indication of the support** (vectors parallel to the gravity, connection to the "ground"); physics. The sitplace is a plane that's parallel to the gravity.

[##*Note from 20/7/2013*: + If it's just one leg, except in some particular conditions (it's glued to the ground; possibly if it starts from the center of gravity), the virtual physics in mind suggests that it would fall; for two legs – it would look unstable. The stool/chair has to seem stable, otherwise one wouldn't attempt to sit, unless she wants to hit the ground with her butt. Notice, that:

* **a plane, in space,** is defined by 3 points – that's why the minimum stable stool or table (they are the same – a plane on a support) have 3 legs
* that's also a suggestion, that "a plane" is a crucial component of the concept of a chair
* also because chairs usually stay on the **floor** or on the **ground**, which usually are **flat**]

> P.S. One interesting thing that is emerging for me from this is that what appears to be a
>question of comparing separate *objects* is vastly more complex than that. We are all thinking
>about chairs, cars, houses and the like as if they were just isolated forms, 2-d,3-d or otherwise.
>In fact, the recognition of those objects is **inseparable from their relationship with human
>beings** - with what human beings do with them. The business of recognizing involves
>comparing and aligning objects like, say chairs, with human forms. We are not consciously
>aware of that part of recognition most of the time, but it is there. There is a lot more here than
>meets the eye (and what meets the eye is hard enough as it is).

Yes, the objects are recognized as different from "background", visually that happens on the bases such as:
-- coordinate adjacency of points/voxels/vertices/planes of the "objects"
-- features continuity
-- experience of parallel/synchronized intentional or non-intentional transformation of the parts of the "object" in contrast to the "background"
-- different speed of translation in the view field
-- size

Stuff like that.

[##*Note from 20/7/2013*: That comparison is not difficult – as noted elsewhere, human-figure can be modeled with about 20 boxes, and comparison and alignment of the objects to human form

may involve just one or a few boxes at a time. For example:
* the size (length, width) of someone's **bottom + thigh**, compared to the size of the seat
* the length of the **part of the leg below the knee,** compared to the **height** of the plane of the seat to the ground plane
* the **length of the back of the chair,** compared to the **length of the back of the human**

]

I'll ask you a few very simple qustions, on your visual IQ:
- Do you draw/paint, how well?
- Have you been into 3D-graphics, computer vision?

All you know is that "it's unsolved problem". I don't think so.

I'd rather call it "unintegrated solution" or "unimplemented solution".

For 3D-reconstruction, study Mark Polleyfeys or photogrammetry.
Other transformations to reduce polygon resolution, voxels etc., and to match images/textures are generally trivial.

IMHO those recognitions will become trivial tasks in a few years.

--Todor "Tosh" Arnaudov
http://research.twenkid.com
http://artificial-mind.blogspot.com

Jim Bromer <jimbromer@gmail.com> Fri, Apr 27, 2012 at 6:55 PM (12)

Todor Arnaudov <twenkid@gmail.com> wrote:

„You mess up abstract concepts with specific concepts, it's not your fault only, I've seen most people does, also in AGI."

Todor,
This is an interesting criticism of some of Mike's reasoning. I hadn't looked at the possibility that **Mike is transmuting an example of a specific into an example of the limitations of an abstraction, but I think you are right, that is what is doing.**

The problem is that an abstraction of x-class is a specification (of the abstraction) and it can then be used as a specification - just as you are doing in your discussion with Mike. I think that thinking depends on that ability. And we use a specific as a placeholder for abstractions all the time. Furthermore it is reasonable to take a number of specifics to use as a demonstration of the characteristics of an abstraction. Mike's error is the belief that all of the specifics of an abstraction (like spatial representation and analysis in computer science) can be held to what he remembers from geometry class.
Jim Bromer

Todor Arnaudov <twenkid@gmail.com> Fri, Apr 27, 2012 at 9:56 PM (13)

Hi Jim,

>The problem is that an abstraction of x-class is a specification (of the abstraction) and it can
>then be used as a specification - just as you are doing in your discussion with Mike. I think that

>thinking depends on that ability. And we use a specific as a placeholder for abstractions all the >time. Furthermore it is reasonable to take a number of specifics to use as a demonstration of >the characteristics of an abstraction.

Sure every system and generalizing hierarchy has something which is most specific. The lowest level input data is the most specific possible for an AGI - "raw" sensory inputs, having intensity, type (RGB for vision; pressure, temperature, pain; kind of taste, smell etc.), coordinates within the sensory matrix, temporal coordinates (within a "buffer"), other timestamps, global coordinates within particular coordinate system etc. The intentional data and drives are also part of the most specifics data into a mind.

There's no "more specific than that" - well, there could be said, e.g. a higher resolution description, a molecular description of the system etc., but if it's available to the particular system, it's either direct "raw sensory input data" or must be converted and represented as such.

Also - everything had to be shaped in some frames and limitations, and have some specifics.

In the context of my discussion, one of the basic "specifics" is simplicity/smallest size representations/buffers of spatial data, like 3D models constructed of the lowest number of planes, boxes and curves, and their relation to the vector of gravity and the plane of the ground.

These specifics emerge simply -- they are highly repeatable **(gravity vector and ground plane)**, it's the **first/shortest to compute** and in the most cases different samples (different 3D models) match to those simplest representation, after their geometry is simplified to lower resolution.

That repetition makes those patterns stronger and implies **"Shape Priming"** for them, rather than to the more complex ones, which can harder achieve a high degree of match between one another, need more operations between match, and on the way **already included matches with the simpler (shorter) models.**

When the unusual parts don't match anything else, of course the features which match are taken as determining the object class, such as those of mine: "a plane and support". What doesn't match anything known (expected) is either a non-relevant specifics for the case, or can be taken as a feature for creating another new class.

-- Todor "Tosh" Arnaudov
http://research.twenkid.com
http://artificial-mind.blogspot.com

Abram Demski <abramdemski@gmail.com> **Sat, Apr 28, 2012 at 12:04 AM (14)**

Todor,

Don't forget to check negative cases. Your definition includes any chair, but also includes tables, floors, shelves, and planks of wood lying on the ground.

I prefer a functionalist definition... a chair is something "for sitting on".

--Abram

Todor Arnaudov <twenkid@gmail.com> Sat, Apr 28, 2012 at 1:38 AM (15)

Hi Abram,

>Don't forget to check negative cases. Your definition includes any chair, but also includes >tables, floors, shelves, and planks of wood lying on the ground.

Thanks for the remark! Yes, you're right, but **you can sit on them all and can use them as "chairs" if needed. Many basic visual concepts don't have abrupt boundaries, the boundaries are sometimes set top-down.**

The woods were probably some of the first chairs and tables*, and also the first specially made wooden chairs and tables were topologically the same, with the only visual difference in their size.

*In Swedish "bord" (board) means "table", in English too. Of course there are many senses defined by contexts etc., but I'm talking about very general visual/3D-models.

What exists in reality **is closer to mere 3D-objects, textures, light sources, weights, sizes, relative sizes... There are neither chairs, nor tables, nor floors as integral building blocks.**

The small **difference between chairs and tables** are for example:

-- in perceived relative sizes
-- proportions
-- few additional optional planes
-- records (memories) of different other items on the planes of the objects of the different classes;
-- some repetitive additional planes and "decorations" etc. (but they can also reduce the set of chairs, if included all the time)

There are also physical constraints related to function -- the **material** of which the items are built, i.e. a table with a glass plate.

>I prefer a functionalist definition... a chair is something "for sitting on".

I agree that it plays a role in creating the concept from scratch (and I'm talking about creating concepts from scratch), I mentioned precisely that functional definition, either. [See email 9, regarding: "top-down … mind searches for a chair … ".]

However function and size are not always unambiguous or known from pictures/visuals (I mentioned that the best is to have absolute dimensions in the 3D-reconstruction), and the example I comment on is the pictures.

I also mentioned about top-down decision of "chairs", i.e. sometimes you should define what you want to find, because part of the classification is arbitrary and visually/3D not essential, there are "decorations".

If one's searching for "chairs" (this might imply something to sit on, i.e. must have the prerequisites to sit on + any irrelevant decorations) then **one will see "chairs" in everything that matches the minimal requirements you've set.**

For example I use my chairs as tables also (when there's no available table), and tables as "chairs" - in fact I just use them as "planes", in programming terms, they are different implementations of the abstract class "plane". (The same goes for floors, shelves etc.)

I've used tables for "shelters" also, while I was little enough. Chairs also can be used for shelters, if the chairs are very big relative to the user. And there are chairs (stools) which look just like

tables, there are small toy tables which are like little stools for the bigger kids etc.

In 3D-terms though, it doesn't matter what you use it for, if you remember your environment as a 3D-model, it can be used for whatever is needed when you need it.

...
A bit of Linguistic Reference to Slavic Languages

BTW, I suspect the arbitrary visual difference between "table" and "chair" is the reason in Russian those words to differ in one single vowel: "stol" (table) and "stul" (chair). In English there's a similarly sounding word for chair without a back - "stool".

http://www.stolstul.com.ua/

In Bulgarian "stol" means both chair and stool (you may specify is it with or without back, how many legs etc. with additional attributes).

In one specific context "stol" means also a canteen, e.g. in a school or University, and perhaps comes from the Russian word for a serving (that apparently is the word for a table), a dinner with many meals, like "a three-course dinner".

In Bulgarian and Russian, there's another word - "prestol", one of its meaning is "throne".

Bulgarian/Slavic "Stol" perhaps comes from the verb stoya/stoit, which has many derivatives, in English this is "to stay", in some of the derivatives it's related also to "to stand", "standing".

...

@All

Anyway, I admit that I talk too much, I should rather make this working for real on its own, then present the results.

-- Todor "Tosh" Artnaudov
http://research.twenkid.com
http://artificial-mind.bogspot.com

Generalization - Chairs and Stools

RE: [agi] The Fundamental Misunderstanding in AGI [was Superficiality]

With: Todor Arnaudov, Mike Tintner, Aaron Hosford

6 messages
Todor Arnaudov <twenkid@gmail.com> Fri, Oct 26, 2012 at 3:12 PM
To: agi@listbox.com

Mike: "...The main reason AGI-ers invariably end up dealing with blocks worlds, is because all their systems, logical, maths, geometry, and computing, presuppose a world made of uniform elements – uniform blocks. – uniform numerical, logical and formal/geometrical units. And this permeates the theorising of AGI-ers at every level. They have blocks hammers so they presuppose the world is made of blocks nails..."

Todor: Many are "stiff", that's right, but not all, and you're classifying wrongly some of those that you're citing.

What physics proves is that world *is* made of "blocks", and what biology proves is that your sensory system is built of "blocks" - finite amount of receptors, and a finite amount of muscles, with finite precision. You just have to have resolution that is high enough (it has two dimension, I've already told which) and you'll have the "endless variety".

Let me generalize your profound confusions once again and discuss about the pseudo-general intelligence of yours and of most of the humans...

--- You don't understand levels of generalization and resolution. You see everything as "different", no matter of resolution or abstraction (there's no resolution in your mind).

As of the picture with the "chairs":
http://www.listbox.com/member/archive/post_content.html?post_id=20121023200141:FAEBD39A-1D6D-11E2-AC2AF39E858D62D4&cid=1C2821765FE74D60B9EE0E5A63117AFB@MikePC

Well, these are not chairs, this is an image, where particular one-color (black) spots can't be recognized on a background (it's a contrast, adjacency/continuinity thing). You say that they are "chairs", this is the top-down part of naming things, you say "these are chairs".

1. "Chair" and anything that is generalized (words are generalizations, except the names) has **some "general/common" parts, features, details, "substances"** - that's the **essential features upon which the word was coined.** (Recall my generalization of what a "chair" visually is.)

2. And... many **inessential features**, "**decorations**", **details** - that are present in some samples and lack in others, which may or may not be recorder as fixed samples, "endless variety". You may put anything in this set, that's **simple supervised labeling**, not bottom-up generalization.
I've already discussed this issue - what you and the top-down-and-bottom-up-generalization-deprived persons are doing is **trying to generalize un-generalizable taught (forced) labeled details**, which apparently lack enough common "basic" features in order to be "easily" and unambiguously grouped and distinguished from other "things". They can be distinguished only

in forced labeled groups, "this elephant is a stone, this toy-car is a stone, this octopus is a stone" (say, because they all are gray) etc. - OK, if you're taught so, you may call them all "stones", and your partner will know that you mean either of an elephant, a car or an octopus, but given the word itself she won't be capable to understand which one of them did you mean, neither she would be able to make inferences about the qualities of any of the items, based on qualities of any other, except their colour which is the same. The "physical" semantic value for someone lacking that particular randomly selected set is close to zero.

That's similar case to what small children often do - **generalizing using inessential features that just has taken their attention**. So you fail to generalize those things, because it's impossible and they are not general, and then you conclude that generalization is impossible, there are no patterns, etc. The fact is that you don't understand what's the pattern, you see a bunch of random details.

--- Mike, you don't have a clue about maths, as others have pointed out...

I've already tried telling, then gave up continuing, but in fact maths is happily dealing with your "blobs" and "patchworks" for centuries, they are called **contours, curves, curvilinear integrals, integrals, functions;** generally your **"endless variety forms" graphically are curves, parametric curves and parametric surfaces, any of them and the families of them can be defined mathematically and they are defined, and all of them can be grouped by some mathematically measurable feature, and by altering the values of the parameters all instances will be produced.**

Such features are:

- closed or open
- number of inclusions (groups, subgroups))
- length of the curve (curvilinear integral)
- area or volume (double or triple integral). Is it continuous or has interruptions, what are the slopes here or there, etc.
- relative length of the curve to the radius, to different radii, to area etc.
- angles between segments having particular features of the above

Colours and their distribution can be treated the same way, differences in the contrast are dealt with differential equations.

Besides, those on your pictures are simple functions, thus simply recognizable, but you don't know what a function is.

-- You talk about creativity, your friend Detusch talk about chess, how creative a human player is, blah-blah...

I suppose you both are artistically-deprived (if not disabled) - creativity in arts is in fact the same like in chess, it's just the average people who are creativity/artistically "disabled" in the part of producing it themselves.

They cannot understand how art is created, they don't understand and don't manage the rules of those other "games" - those rules go beyond their cognitive capacities. Music, drawing/painting, photography (composition, light, contrast), dance, creative writing, sculpture, acting, juggling...

I**t's all so trivial, and mathematically elementary.** What do you people don't understand about art, creativity and about the creation of pieces of art or new inventions? **Please, consult yourself with lectures, materials, papers, and then you should understand all of it.** If you fail, it's your brain fault and your pseudo-general intelligence.

In fact , the reason why intelligence appears so "magical" in general lays is in the pathetic

cognitive capacities of the average humans. Intelligence is trivial, but humans are not smart enough. By the way, **the general intelligence of humans is a myth**, more specifically **what's called "general" is not that "general"** and **is not symmetrically general.**

As I've discussed here also, human cognition is asymmetric, people are **largely passive recognizers**, **not creators/producers,** and their failure to understand and to produce super elementary things, in addition to their associations with emotions, make those appear as extraordinary.

One "generalist" here may claim that it's "a trade-off between precision and generality", but I'll say that this is a bullshit, I challenge his intelligence "generality".

What "precision" you need for drawing, or for playing a piano?

It's elementary perspective transformation, layers, and trajectories. In the case of music - it's hitting super elementary sets of adjacent or otherwise-related keys, building-up scales, chords, arpeggios, melodies. People are amazed how did Beethoven composed music while deaf? Well, even if you were deaf, and you were given a piano keyboard, what could you do with it? Eat it? Melt it? Cut it? No, you could press the keys at different velocity, hold it for different length of time, at different coordinates, in different successions, you could repeat those, you can use all of your fingers at once, do it in different scales, combine, variate, transpose, change the directions etc.

That's the "magic" of musical composition and it's mathematically elementary. The great composers are just great mathematicians (in average humans scale) who have conscious access to the musical data (pitches, times, ticks, their relations) and much better working memory than their audience. If physical motion is blamed, such as "fast coordinated motion requires a lot of practice" - well, do it slowly, can you?... But no, average people, including "generalists" can not deal with music or with drawing, or dance. What's "so precise" here, what amount of cognition is needed for just a bunch of a small set of possible motions done in particular way and synchronization with music or other "clocks"?

So what's "general" about average human mind or your own mind's intelligence, if you can't generalize and understand such obvious things in their native domain? (Includes maths, soft sciences etc., you should deal with ALL with the same mind, **otherwise you're a "pseudo-GI"**, or a "partial GI")

I will conclude: It's not that GI is complex, it's that average humans' minds are not "complex" and general enough, the AGI will be intrinsically more general than a human mind, it will be in fact more of a "general intelligence" than "human level" (average human, poor "talents", or a few specific talents).

So average people or ones who don't understand art or creativity may accept the simplicity of it only when you slam a thinking machine, that will outperform them, right into their faces. That's what any of us who wants to create real GI should do...

---> Todor "Tosh" Arnaudov <---
-- Twenkid Research: http://research.twenkid.com
-- Self-Improving General Intelligence Conference: http://artificial-mind.blogspot.com/2012/07/news-sigi-2012-1-firstsigi-agi.html
-- Todor Arnaudov's Researches Blog: http://artificial-mind.blogspot.com

---

Todor Arnaudov <twenkid@gmail.com> Fri, Oct 26, 2012 at 10:13 PM
To: agi@listbox.com

@Aaron,
>Just give up, Todor. It doesn't matter if you're right or wrong. He's taken a stance that won't let him think he's wrong, even when he is. He'llre-
>interpret what you've said to something else so he can pick it apart. He'll change the subject when he can't come up with a good argument. He'll
>directly contradict the truth by denying you've done/said something you just did, or by relabeling it as something entirely different. He doesn't want
>to find the truth; he wants to convince you he's already found it. He doesn't realize that there's a trade-off between feeling right and being right -- that
>the humility to recognize your own failures and shortcomings as such is the very thing that makes you able to overcome them. It's more important to
>him to look right than to actually be it.

Thanks for the advice... You are right, and I always give up, then give another shot. :X I guess I'll give up in this
iteration in a bit... :)

Mike:>Wherever you look in the natural world, you DONâ€™T see uniform blocks â€“ you see groups of irregular, individual BLOBS â€“ indeed PATCHWORKS of blobs â€“ >rocks, earth, cells, tissues, living bodies, theyâ€™re all made of blobs, not blocks..

Todor: Have you ever heard of "particle physics", atoms, protons, neutrons, electrons, quantum physics, quants. A "blocks" means a "building blocks". Yes, in nature there are more of spheres than rectangles or cubes, it's a simpler form, requiring less variables to define, but it's a "building block either".

The first *known* hypothesis for the atoms is made thousands of years ago, it's possible that people thought that tens of thousands years ago.

What physics discovers is that world is ultimately made exactly of uniform "blocks", having absolutely identical properties.

Mike:>Science always makes mathematically simplified models of the world. What youâ€™re arguing is equivalent to
saying Newtonâ€™s calculus proved that curves are made
>of rectangles. Theyâ€™re not â€“ thatâ€™s a useful and brilliant formal simplification â€“ in fact, translation.

Todor: BTW, you're repeating a philosopher I've been arguing a decade ago, he was blaming science for "simplifying things" - well if you can understand things and can predict their behavior, you "simplify" them. If you cannot understand, you take it as a whole and just copy it. Yet the ones who don't understand it (and can't simplify it in a workable fashion) cannot really understand the "simple things".

Please design for me a simple digital computer, say equivalent of an Apple][, or not that complex, let it be like Whirlwind. Those kind of "simplified" concepts are simple only in theory, and if taken out of their complex and real application.

Yes, "curves" and integrals are simplifications, the real world is made of bosons and fermions, of 10E+3894938493 particles - so how should you process it like that? How do you process electrons or quarks, and did your retina or your brain know initially what an electron is?

"Curves" are just images on your retina, or sequences of neuronal activities when touching objects etc., and maths allows if you know some measures of those curves (in terms of the sensory matrices) to predict (compute) some other measures, which are not directly observable by the regular sensory matrix.

Mike: >Rational technology INVENTED uniform blocks â€“ it was one of the greatest, most imaginative inventions in history â€“ and you â€œblock-headsâ€ think they
>*discovered* them - bricks, for example, and perfect rectangles, you think, *always* existed.
Nonsense.

Todor: Your notion of "imaginative" is pathetic, perhaps you think this is "the most imaginative art" are the rectangles of

Malevic:http://en.wikipedia.org/wiki/File:Malevich.black-square.jpg

Uniform Blocks

The uniform blocks, especially triangles or rectangles are the less imaginative (the easiest, the simplest - I've explained this before), and **Universe is built of uniform blocks**, nature at larger scale, seen with naked eye either:

- insects (see a bunch of ants from the above),
- their eyes;
- bees's cells
- nails
- hair
- grass
- sand (!)
- leaves
- fingers
- eyes
- trees in a forest
- animals from the same species - see a flight of sparrows from a distance; or
- a pack of wolves; or look at
- the stars or the planets; humans either.

- **Connect 4 points and you get a "rectangle"**,
- Try to build a house with the fewest possible number of walls/parts and you get first a triangle, then
- Make a compromise and use 4, which will give you twice the area for a small overhead - and there you go, that's a rectangle/tetragon.
- **Why trying to do it with the fewest number of elements?** Because it's a hard work, you get tired, have to hunt, to walk around in a search for food etc. That's why one would search for the simplest solution at the moment.

Is it clear now how it can be "invented" from scratch, how that "new element" can be created... Of course, for humans who cannot generalize it, that's not rectangle, these are just "trees", they are "not uniform", "don't have exact dimensions" (in pre-stone age - right) etc.

You reason so, because you cannot get the meaning of "generalization" and "resolution". For example, let's see "the greatest invention" - the wheel.

Wheel was invented? No - it is obvious

Well, who said it's so great? It's absolutely trivial!

Have you ever looked at the sky? In the stone age there wasn't TV or light bulbs, so **you were supposed to see the moon and learn what a circle is from the very early age.**

In fact you learned even without the moon - **have you seen your mother's eyes?** They are **circular**, they have several circles, they are even **spherical**. Humans have probably seen **the eyes of dead mammals** either.

- **The moon is a "wheel"**,
- the eyes are wheels and **they roll**.
- The **trees** are also circular and **can roll**.

  There's nothing ingenious in "inventing" the wheel.

The problem is that modern people usually have a hard time thinking in terms of the context of the inventions - first they don't realize all those examples of wheels, and second they don't realize that **the wheels require roads.**

How would you use a wheel for traveling, if you're surrounded by a dense forest and all around is muddy?

For the wheel to make sense, roads are needed, or a **dry place with flat areas** where the wheels/cut stems of trees could roll. In case of trees – you **must have technology and energy to cut trees that are big enough, and to slice or "delve" them to make wheels,** or to have
technology to dry and bend wood and connect pieces in order to make discus etc., and also have the energy to cut enough trees to make a road etc.

I.e. the wheel itself is not the problem, the other technological difficulties are. Also you must have horses or other strong animals (they had to be domesticated first) to drive such heavy clumsy and probably not quite circular wheels (having a lot of friction), otherwise they are useless for transportation - not surprising that the first wheels were used for making pottery(?).

Mike>The problem of AGI is always â€“ if weâ€™re talking formally, quasi-mathematically â€“ to deal with blobs and patchworks of blobs.
>But obviously, this kind of discussion is too high-falutinâ€™ without examples, and specific analyses.

Todor: A bullshit, I've analyzed your "blobs" too, and give more than explicit examples and analyses.

Mike:>The general point is: you have not and will not show any uniform blocks to underlie those chairs, or the concept of
CHAIR â€“ or **any concept period**.

Todor: In fact I already have shown the uniform blocks of the concept of "chair", regarding your last "blobs" - sure - those are not chairs - "coma". They are "chairs" only in your head, because you say they are chairs or because you have copied them from images labeled as "chairs". In fact it's just 2D black-white image - "exclamation sign", you can label any thing as a "chair", **in fact it's just an image, in this case the rest is arbitrary label.**

The "blobs" on that particular image can be represented exactly in the resolution of this image as very simple functions - that means if using those functions the same looking "blobs" can be redrawn at another place, and by adjusting or transforming their free variables they can be drawn as bigger, smaller, inclined, rotated etc. If one see those and the other "blobs", and asked, he'll recognize them as similar.

As of the real (3D) chairs - I take the challenge, my thinking machine is supposed to categorize them and give them a name without supervision, also to design any "infinitely varied blah-blah

never ending" new kinds of chairs, I even have explained how it works. However I just note for the ones who do get generalizing - they will seem as "radically new, novel elements, blah-blah" for ones that don't understand generalization. In essence, their "chair-defining" properties would be the same, otherwise apparently they wouldn't be recognizable as chairs, except in - as mentined - arbitrary labeled sets, which are randomly labeled.

"Concepts" live in your head and in mind, in the real world yes, there are no "concepts", but mind cannot operate mentally with the real world, it can operate only with concepts derived from its sensory matrices space, and the sensory matrices are finite, "rectangular", "circular", "uniform" etc.

Much more objective elements are "the particles" and measures based on their properties; if you prefer "wave functions", "number of hydrogen atoms diameters" etc. but unfortunately for your point, these measures seem to be exactly uniform, regular etc.

Besides that brain or any system except the original Universe cannot process them at their original resolution. The experience have shown - brain or a machine don't have to model them at their original resolution, a TV with a low resolution or a video of 320x240 or half this is enough to see a lot. The mathematical physical models work ever finer (and their resolution grows, though), i.e. they produce correct predictions.

Mike:>If anyone wants to talk specifics and analyse forms as I did, Iâ€™m delighted. Vague generalisations alone don't cut it

You did what??? Yes, "vague" generalization that you can't get don't cut for you, it will "cut" your bullshit when we do create the thinking "terminators" to finish your confusions. (Relax, just the confusions.)

Todor Arnaudov <twenkid@gmail.com> Mon, Oct 29, 2012 at 3:42 PM
To: agi@listbox.com

Aaron, you're right that we should stop and go working, than explaining. :X But...

Mike> Challenged to explain what are the uniform blocks - or elements – or semantic net units –
> of chair, Todor, PM and
 >Aaron, have all responded in the same “essential” way.
>“My machine will do it”
>“My program will do it – it’ll learn how”.
>How will the machine/program do it? What are the blocks they will identify?
>No reply.

Todor: LOL, no reply, really?... :)) I explained you about chairs months ago. This is how, in a few simple steps:
- Identify the gravity force,its direction
- Identify a support plane/surface
- Identify a plane that's perpendicular to the gravity vector and has a support
- Size must be appropriate for sitting (area/length of the planes of the ass, of the chair)
(Some additional)
- The center of gravity must be low enough so that you don't fall after sitting
- The sitting plane has to be reachable, and must be big (or small) enough

There you go.
If anything recognizable as a "shape" (3D reconstruction through vision/projection/reprojection/ RGB-D cameras, laser scan - whatever) has those features, it's a "stool", more precisely i.e.:

The back is just yet another plane (or more generally a surface), that is connected to the sitting plane.

What's a surface? That's a connected set of "voxels". How do you know it's connected? Well, visually, there's low contrast between the pixels and then the voxels. The low contrast between them, in contrast to the high contrast between their borders and the outer part of the scene, makes mind to cluster them as entities. Namely their "borders" are the coordinates of highest contrast.

>.Instead: How could I possibly be so stupid as not to understand what they have said? So >cussed? So have-to-beright?

Todor: Bullshit, endless explanations go to you, precise and detailed, but you don't get it.

>What are the blocks again?

Todor: I told you many times, and your notion of "blocks" is super primitive.

>What did you say were the blocks?
>God this man is impossible...

Todor: I told you months ago, you don't remember the previous messages.

Initially pixels, then voxels, planes, surfaces and forces and they correlations, both in a global space or in a local-internal space between a set of pixels, voxels, planes, surfaces.
If you can't draw a picture out of this explanations, it's your imagination's fault.

>It's quite possible that never in history has there been a would-be (but not actually) creative
>field that has been populated by such a phenomenal percentage of bullshitters. It's not just you
>guys – it runs through the field, and extends to the highest echelons present and past.

Todor: Mike, you think you're creative? You believe a rectangle is creative, or your random "blobs" that are described by 5 bits-long algorithms.

>If there are blocks, or elements, to a chair – or any concept, period - you can draw them.
>You ain't drawing them.
>To paraphrase Feynman: if you can't explain it simply, you almost certainly don't understand it.

Todor: Yes, I can draw them - comma, I'm an artist, too, I understand perspective, 3D-reconstruction, light. What about you?
(...)

I don't draw those planes, because writing is clear enough and I'm lazy doing it, I assume the people reading have enough of visual imagination to imagine a few planes or surfaces, but it seems it's not that easy.

>When you try and draw the blocks, if you ever try and draw them, you'll realise there are none.
>(And you might just begin to understand the formal beauty of blobs and patchworks).

Todor: You should go study some Calculus, Linear Algebra, Analytic Geometry, Differential Geometry and Topology - you'll find the answers - I told you what your "blobs" and "patchworks" are - curvilinear integrals, curves, surfaces, double and triple integrals etc.

The "conceptual blocks" are not static drawings, they are correlations and patterns, sometimes hierarchical, sometimes not, and sets of transformations etc. However that is maths...

* Todor "Tosh" Arnaudov *
* Twenkid Research: http://research.twenkid.com
* Self-Improving General Intelligence Conference: http://artificial-mind.blogspot.com/2012/07/news-sigi-2012-1-firstsigi-agi.html
* Todor Arnaudov's Researches Blog: http://artificial-mind.blogspot.com
[Quoted text hidden]
-
Todor Arnaudov <twenkid@gmail.com> Sun, Nov 4, 2012 at 1:23 AM
To: agi@listbox.com
"Todor:

- Identify the gravity force, their direction
- Identify a support plane/surface
- Identify a plane that's perpendicular to the gravity vector and and has a support
- Size must be appropriate for sitting (area/length of the planes of the ass, of the chair)
(Some additional)
- The center of gravity must be low enough so you don't fall after sitting
- The sitting plane has to be reachable, and must be big (or small) enough
There you go."

Mike:> Well, at least you're trying.

Todor: See also the old explanations, I'm lazy to search and copy.

>This isn't really coherent – it sounds like you're saying that all these chairs must have a seat/"sitting >plane"/"support plane" to classify as chairs.
>But what does a support plane or sitting plane look like? Draw it. There are a lot of support planes >on these objects, just as gravity is acting on them at many different points – and there isn't one >concentrated direction of "the gravity force."

Todor: "A plane" is a mathematical concept. It looks what a "plane" look, it's not a drawing, it's a correlation of points in space, it can be in voxels, a volumetric integral (if the thickness is taken), or surface integral which if seen in low resolution is just a plane. If it doesn't turn into two planes or if there's no plane after resolution reduction - that would be something else.

There are not a lot of support planes, this is the lowest resolution representation, and there is a concentrated direction - it's an "averaged vector" and goes through the center of mass, which usually is through the center of the plane that is perpendicular to the "ground", the ground is found by the direction objects fall if there's no support etc. (Center of mass is what allows drawing realistically human figure in action, and it is one of the ways to anticipate motion in picture and that somebody is going to fall - when she's out of balance, which is determined easily for the human body, apparently by our experience of watching humans moving and falling).

>If you mean that such a plane is any that a bum can sit on – what does the bum (or human-figure-
>with-bum) you are going to apply to these drawings look like, and where are you going to apply it?
>How do you know where to apply it – where and how to sit on these chairs?.
>And if your sole criterion of a chair is a seat/sitting plane, how are you going to distinguish chairs >from swings?
Todor: I have already written about that and answered some of those question.
Yes, there are many criteria, including the materials, here I'm talking about the visuo-spatial static domain (voxels/contours/planes) and gravity ("statics" in physical terms). There are many specific details, which are not essential. A "swing" is obvious - it can swing and naturally swing ("swinging" is the trajectories which are called so) - it's not static, it's dynamical, i.e. the coordinates of those planes are supposed to change. Yes - they can be changed for chairs with wheels, and one can swing on any chair (but may fall back, too), but the directions are different. In visuo-spatial low resolution domain, they are similar, the details, materials, static or dynamic, the axis where it rotates or swings, the location of the chair/swing in global coordinates (the context) etc. etc. of course add to the definitions, and all can be classified in sub-classes for specific types of chairs, swings or whatever.

I am talking about general cases, in low resolution, otherwise there couldn't be words, but just exact pictures, not concepts. I've already discussed precisely this issue for chairs and buildings.

>Your method doesn't add up to a coherent form – a "uniform block" - that all these chairs have in >common –

It's the opposite.

> or a coherent method of sitting on them, and identifying them that way.

Todor: The "method" is to compare the size/surface of the bottom to the surface chair, the height to the ground, the strength of the support, of the back etc., that defines whether you can sit and will you fall. Those comparisons are literary comparisons: matching the size/area/height, numbers. There are many ways of "sitting" etc., but "stool" or "chair" are *general*, these are words - their purpose is to tell as much as possible for the context with as little data as possible.

>P.S. And what happens if you're confronted with a chair upside down, or whose seat has been >smashed? We can still recognize a chair if the seat is missing. How?

Todor: Upside-down is an invariant transform of a 3D object, it's a rotation and/or translation, an object transformed that way (the system of simplified planes) its the same in visuo-spatial voxel terms. (Also, if you ask a little child what's this and show her a strange chair that's upside-down, she may not tell you a "chair"). And since you understand classical physics (everybody does in one extent or another) and those transforms, you can see how to transform the chair (if laid on the ground) so that you may sit on it as usual. For average chairs/stools, there are some frequent additional features, such as 4 legs, or legs, or a pillar with wheels, or a
lever for adjusting the height etc. also there is context - where this "thing" is - in front of a desk, a table, ... (another plane) where people are seen to sit. Also recognition is possible, because parts of chairs are different enough from parts of other objects (otherwise it'd be ambibuous).

What does "smashed" mean - to what extent? If there are other parts that match, they are still recognized, the plane is still there (if the legs still stand, or if the chair is in pieces), the legs are still legs, the back is still back (a plane/surface), the sizes are still the same and fit an average or someone's bottom, there are wheels on a pillar with appropriate size etc

>Your method doesn't add up to a coherent form – a "uniform block" - that all these chairs have in >common – or a coherent method of sitting on them, and identifying them that way.

Todor: Mere text explanations are not a formal working system, the details are gotten from the data, and it's faster than explaining it detail by detail. Whatever, there's one way to prove and show what I'm talking about, it's not by fruitless discussions and explanations.

--
* Todor "Tosh" Arnaudov *
-- Twenkid Research: http://research.twenkid.com
-- Self-Improving General Intelligence Conference: http://artificial-mind.blogspot.com/2012/07/news-sigi-2012-1-firstsigi-agi.html
-- Todor Arnaudov's Researches Blog: http://artificial-mind.blogspot.com

---

Todor Arnaudov <twenkid@gmail.com> Sun, Nov 4, 2012 at 11:17 PM
To: agi@listbox.com

Mike, these are not chairs [see the annotated image], and my method requires 3D - the *planes* require 3D to be defined. In 2D there are only lines, (closed) curves, (closed) contours (you call them "blobs") etc., and this is a 2D picture, or a set of 2D "objects" (identified/separated by proximity/and low/high contrast), I already told that.

The visuo-spatial concept for "chair" is derived from 3D, where the observer can play with the coordinates and orientation of the objects and the "camera"; in a space with gravity, with known (perceivable) dimensions/sizes/distances/strengths of the materials etc., and known sizes of the agents/their parts that would use those chairs, and to which the chairs should match, such as:
- average length of the thigh
- average length of the leg from the knee to the foot
- average length of the back (of a human)
We reconstruct 3D from 2D by reverse-perspective and reverse light transforms, and because the G-force and center of mass recognition (that's one way to recognize intuitively where's "down"), and already seen 3D things which project in 2D like given 2D images etc. (some matches are recognized directly, like the legs, some ranges of proportions).
The recognized archetypal cases, their legs (gravity), relative dimensions etc. also suggest that the strange ones "might be chairs".

However **3D from 2D is an ambiguous operation**, it cannot always work correctly (the reason for some of the optical illusions), and that goes even for shaded or color images.

Obviously generalizing a class of 3D-objects from a poor 1-bit 2D silhouette projection from one point of view is vastly ambiguous. I already told about that - using vague digested features that cannot be generalized as expected, because **the general in the originals is swept away in this representation.**
...
Also, the recognition of given objects as "chairs" (and not instances of another class) is due to the fact that the other known things/samples are *different* enough.

In simplified representations and with a whole lost dimension those differences get reduced, and different classes now look the same (that's one aspect of resolution reduction I talk about).

What if I see it in 2D, as it is? You and the context (the most "archetypal" cases) suggests that perhaps the rest are also part of the same class – **it is you the one who suggests that**, I also already explained that, but... If any of these pictures is shown separately out of the context and you ask a random guy to tell you what he sees in one word - I don't think he would always say a "chair"...
[*##Note from 20/7/2013*: Experiment to test this! A survey - but I need more images from different classes. :)]

Let's see the figure that is 3-rd on the right, second row from the bottom - nobody would say this is a chair if seen that way in 2D in 1-bit color separately from the context. This is rather **a satellite receiver, a deep dish, a cone with a "halfmoon" on it**. The one on the left looks like a head put on a pillar - it has an ear, a nose, a chin. The one next on the left
looks like a barbeque or like something to put a baby in.

http://pristine-home.blogspot.com/2012/04/barbecue-time.html
http://www.diytrade.com/china/pd/6599510/Electrical_Baby_Swing_BSE900.html

The one below looks like it's made from balloons.

There's a "magnifying glass", above the "head on a pillar":
http://www.mallcarts.com/vdir/vitem/AM-AC099A

Fig. Silhouettes of chairs – tagged by Todor as what they really look like

There's also **another satellite dish** - the right-most in the first top row.

The figure that is two positions right looks like **an (alien) insect-cyclops**.

The fourth on the bottom row is like a face of a cartoon villain who's spitting something - a spider's web or a puke?

The right-most on the top-row is like an **alien ray-gun**, **an antenna**, **a drill, a brush** - **not a chair**, because **from that POV in 2D you can't tell whether the short lines are flat, deep and rectangular, or circular.** The one on the right looks like the Alien, or like two burning torches, put on rest "on the thing in the middle" - possibly an alien holding those torches.

Etc.etc. - see more suggestions on the picture.

If you get:

- a 3D view
- the size
- the materials of which those are made
- their strengths
- their transparency
- + reflectivity
- texture etc.

Then you will get more clues to decide are those really "chairs" or are the other ambiguous objects in the 2D-silhouettes, toy chairs, small chairs, big chairs, designer chairs or sculptures.

The method with planes applies for their most general visuo-spatial structure, and I repeat: **it's for abstractions, for words. In order to specify everything, many words and lower level data are required, and the abstract classes get overlapped**, I've discussed that and the raw data in the initial "chair thread".
…

>Now show us how any of your methods apply - how they can identify common blocks/elements of these chairs/ chair images.
>You don't necessarily have to do the whole chair – just show us how your methods will *begin* to apply to parts.

Todor: You're challenging a method which talks about planes and 3D, using a 2D picture having one plane with everything laying on it. **In one plane you cannot fit a plane, a perpendicular vector and another perpendicular plane, parallel to the vector.**

>You are doing here in micro what AGI systembuilders incl. B&B have done in macro – spent a lot of time detailing a system, without ever taking the trouble to see whether the method/system actually applies to the problem – when in fact you should all
>have *begun* with the problem.

Todor: Sorry, a bullshit, I have visual imagination (tautology), it seems not everybody does. **The visual (geometrical) generalization is essentially linear algebra, topology, calculus, and it's obvious.** It also emerges from the raw data **by simple matching and resolution variations.** (...)

>All your time is wasted unless you can show your efforts really apply to the chairs - and can actually identify similar elements in different chairs.

Todor: I can identify similar planes in a 2D picture, I can identify similar curves, similar sizes of them, similar topology, similar curvature etc. That's **calculus and topology** (e.g. **graph theory**). I already wrote about those fields of maths.

[##*Note from 21/7/2013*: M.T. fails to get the essential "similar element" to some abstract concepts, when all the rest is "different".
]

>I don't think – and, I would say, the great majority of philosophers and psychologists – don't think you will.

Todor: I'm probably partially repeating myself, but most philosophers and psychologists (or I assume yourself, too) don't have 1/15 or 1/20 of my expertise or my talents, or of the talents of other polymaths and universal artists... And yet those "just philosophers" or "just psychologists" believe that they do understand general intelligence or creativity better than them/us...

I've done original philosophy in real time, on the spot, as a teenager, "abstract" ethical and whatever problems are trivial to me - they don't understand in formal terms serious mathematics or physics even in their 70-ies, nor can deal with programming/CS/computer graphics (a close friend of the above, which have many sub-branches) - for most people programming is absolutely impossible to grasp even for elementary tasks, the more complex tasks require even more complex mathematics and too big a working-memory.

On the other hand for ones who deal with it, or for me. it's "like ringing the bell" right from when I first started – trivial. The "average" ones may know some neuroscience, but not for sure, and they cannot make the associations to mathematics and CS, they're "mathematically blind". Also perhaps blind in terms of physics. They probably lack serious visual processing capabilities, like imagination of an artist and a movie maker - to see movies of the things and how they can be transformed or built in a sequence, and to see the outcome of many transformations immediately, without having to do them in reality.

They probably lack the visual and physical processing of an artist or engineer in terms of easily playing with perspective, 3D, mechanics (support, center of mass, balance, statics, kinematics, dynamics, projection, dimensions/sizes...).

If so - yeah, probably they cannot see what I'm talking about, everything shows that they lack a lot of cognitive faculties that I have.

* Todor "Tosh" Arnaudov *

-- Twenkid Research: http://research.twenkid.com
-- Self-Improving General Intelligence Conference: http://artificial-mind.blogspot.com/2012/07/news-sigi-2012-1-firstsigi-agi.html
-- Todor Arnaudov's Researches Blog: http://artificial-mind.blogspot.com

[Quoted text hidden]
chairs-2d-other-tags-captions-twenkid.jpg
125K

Todor Arnaudov <twenkid@gmail.com> Tue, Nov 6, 2012 at 1:31 AM
To: agi@listbox.com

Mike>But unless I missed it, your proposals don't call for simulated interaction of the robot's/viewer's body with the objects >viewed – such simulation is certainly necessary for us to recognise those chairs- and there is massive evidence for it.
> (Is anybody in AGI or robotics attempting such simulation?)

Todor: Mike, I've indicated my school of thought a thousand of times here and in what I agree with you, it's mainly sensori-motor generalizing hierarchies, it is biologically inspired, but improved. **My machine will start embodied, but of course in virtual worlds.** I don't have $ for robots, and IMHO the robots are the worse way to do it.

Once a working AGI is available, it's supposed to manage any robot body through an interface, like a human can use controllers or exoskeletons for controlling another body.

I'm tired of those endless explanations and no understanding of what I'm talking about, though. I'll tell you something I've already told again - yes, **there is an endless diversity in nature, but it's not essential for intelligence, the structure and predictability is what allows intelligence.** (...)
Well, I went into too much of explanations, but I'll rather put them elsewhere in a publication or so - sorry, I have to shut my mouth up and start implementing. :X

Then I will show what my thinking machine prototype can do. I believe I might have something by the next year, if I have the opportunity and mood to work hard on the implementation.

.....Todor "Tosh" Arnaudov.....
-- Twenkid Research: http://research.twenkid.com
-- Self-Improving General Intelligence Conference: http://artificial-mind.blogspot.com/2012/07/news-sigi-2012-1-firstsigi-agi.html
-- Todor Arnaudov's Researches Blog: http://artificial-mind.blogspot.com

4. Caricatures

From: "Mike Tintner" <tintner@blueyonder.co.uk>
To: <agi@listbox.com>
Subject: ..and Visual Individual Recognition
Date: Sun, 13 May 2012 13:41:54 +0100

Another key thing is that recognition is not merely a matter of recognizing a generic body, - a "human","man", "well-dressed man" etc, - but an individual. Lyndon B, / George W.. etc

And the strongest clue that recognition is not geometric or anything to do with patterns, lies in the human ability to recognize caricatures of an individual:

http://2.bp.blogspot.com/_HR3JtqIfTiE/SpXTjQJFNTI/AAAAAAAABDg/Z98CBoODYCA/s400/Lyndon+Johnson.jpg
http://www.history.com/images/media/slideshow/lyndon-b-johnson/lyndon-b-johnson-listening.jpg

http://4.bp.blogspot.com/_y2Jh2toqU7E/TMgUZmjl65I/AAAAAAAAACU/ZC1eC9uaY8U/s1600/george_w_bush_439195.jpg
http://i.dailymail.co.uk/i/pix/2010/03/09/article-1256558-021BA7C7000005DC-283_224x423.jpg

Caricatures are not patterned, geometrically consistent distortions/transformations, but rather patchwork, crazy distortions/transformations. Radically-different-while-still-somewhat-similar-forms/features and outlines are substituted (as in series of patchworks), None of the current AI techniques or indeed any purely-analytic/non-holistic approach to recognition can work.

...

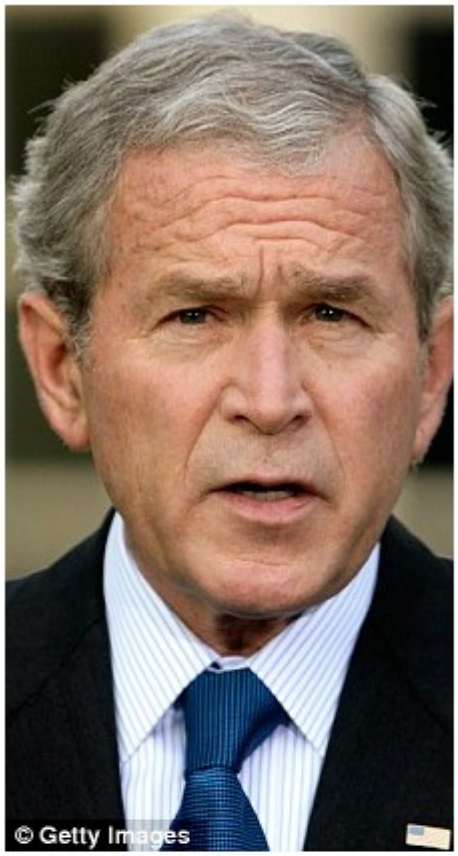
© Getty Images


WILL LIE
FOR WAR

Re: [agi] ..and Visual Individual Recognition
Todor Arnaudov <twenkid@gmail.com>
5/13/12 to agi
From: Todor Arnaudov <twenkid@gmail.com>
To: agi@listbox.com
Subject: Re: [agi] ..and Visual Individual Recognition
Date: Sun, 13 May 2012 22:46:41 +0200

Hi Mike,

I'm sorry to "show off" my visual processing skills again, but I think you (and many others) just fail to notice the patterns, geometric models etc.

I see the following at a glance:

-- **The shape of the eyes**, more precisely **the direction of the line forming the eyelids**. The extreme angle is what defines his eyes' specifics, it's more extreme than average, and anything more extreme than average is just "more extreme than average".

-- **The distance between the eyes** -- it's smaller than in an average face. Anything smaller than average is "smaller than average", so that's why when depicted even shorter it's still "correct".

-- **The size of the ears and their shape** (among the samples of human ears - there are not too many distinctive basic types). His eyes are bigger than average.

-- **The curves of his nose** - among a set of noses (you know - Roman, straight, narrow, wide (black-race), short ... -- they are just a few, and they are a few, because the essential features are a few, i.e. comparison between average values of the "physical features" which are just absolute and relative: coordinates, distances, angles, sizes, color intensity/saturation etc.)

-- **The eyebrows type** (among the samples of human eyebrows) -- wider than average. (Anything wider than average is just "wider than average" - to repeat again.

-- **The color of the hair** (white)

-- **The upper lip** -- it's thin in the original (thinner than average, thinner than his lower lip), here it's completely missing

-- **The shape of the jaw**, i.e. the angles between the chin and the cheeks are of the same extent as the original (they are a type among all types -- square, round, triangle, round)

Those curves, shapes, lines -- features -- can be gotten by emerging contrast and search-for-matches tests, and resolution reduction up to points when only a few distinctive classes are gotten - first as a few distinctive classes of features (eyes, eyeballs, eyebrows, eyelids, nose, lip, jaw-shape, ears, hair, chin...)

The differences to the average features are compared - **the features are a few**, not millions, I mentioned most of them above - and also the actual matches to the samples (records) of those features, which are of **low resolution - traceable.**

"Ear" is the "shape" which starts from there-, ends there, crosses this and that, etc.
"Ear" is the "shape" which is different enough from this and that (eye, nose,

-- **The highest match wins**, sometimes there are ambiguities (look-alikes). The price of recognizing those "fluid shapes" is actually low resolution of comparison, based on those

extreme/to-average features, which allow for endless number of variants, thus -- many different faces are recognized as "the same", which might be wrong in some worlds, when there are clones.

-- There might be a world, where the individuals have visually the same faces, with differences only in the texture only in some very specific places, which requires close look to check. (In computer games it's often like that). In those cases, facial recognition will be useless, like it is if the spatial and contrast resolution is too low, e.g. 3x3 pixels, 1 bit. :)

-- Also, if this caricature is shown to someone who hasn't seen Bush, but knows somebody with matching set of the features above, she wouldn't recognize Bush, but obviously that guy. The recognition would be "wrong" according to you.

>Caricatures are not patterned, geometrically consistent distortions/transformations, but rather patchwork, crazy >distortions/transformations. Radically-different-while-still-somewhat-similar-forms/features and outlines are substituted (as in >series of patchworks), None of the current AI techniques or indeed any purely-analytic/non-holistic approach to recognition can >work.

Caricatures are the least of "crazy" or "inconsistent", artists know exactly what to exaggerate, in order to keep the face recognizable. That's the least of a magic, it's about being observant. I mentioned some principles I'm seeing above.

It's ironical that in fact the caricatures display how human facial recognition works and what's essential, and that there are "patterns".

As I mentioned, caricatures are based on:

-- **Obligatory** (see why below): Keeping the features of the caricature to match the original (a picture, a 3D-model) at a higher level of abstraction. That is e.g. the difference, compared to an average, i.e. bigger/smaller, it's a kind of comparison of "derivatives" in mathematical terms, are the functions "rising" or "falling". A lower level of abstraction is the specific angles, distances, sizes etc. - eg. the angle of the line of the eyelid. However again, the comparison has some degree of precision.

-- **Allow variations** -- keep the features match in lower level too: Some of the features are kept matching at lower level too, like the angle between the chin and the cheeks.

And all those, - either if derived from a stereo picture with two naked eyes, 3D-reconstructed; or by 3D-reconstructed from a 2D projection from a picture - are then reduced to 2D, one of the dimensions is spared.

The **"exaggeration"** graphical implementation is a trivial concept in computational terms:

-- **Increasing/decreasing angles** (the angle of the eyelids in the Bush' picture)
-- **Increasing/decreasing distances** (the distance between the inner corners of the eyes in the Bush' picture)
-- **Some of the distances are radii of curves, circles, ellipses etc.**
-- **Increasing/decreasing relative and absolute sizes** (ears in the Bush' picture)
-- **Increasing/decreasing color/saturation intensities** (e.g. a slightly dark-skinned person, rendered completely black, or a pale persons painted completely white)

The "exaggeration" is also trivially measured by computational techniques.

...

If the requirements above are not met, thus the changes are in the wrong direction in many features, the work would **first look like an "ugly/wrong/bad" caricature, and after a threshold of mismatches is passed - as another unknown or known person**, if the

differences are too much and the matches to the other person grow (see "facial composits", "photofits").

In some cases, when it's still "recognizable" it would be because of trivial and dumb features which are "easy" because of the triviality/popularity of the topic at the moment, and the lack of competing similar cases.

E.g. in a set of 10000 persons which are selected to look very similar, recognition of one of them from their caricatures will be very hard or impossible - their caricatures would just look "the same" after their specific exaggerated features are exaggerated even more.

One of the steps for general recognition is the emergent algorithm to manage to reconstruct 3D objects and textures from 2D-input (mono or stereo) in resolution that is high enough, in order to get "invariant" canonical representations and then use them in any lower dimension and resolution or transformation.

Cheers,
- Todor "Tosh" Arnaudov -

First Conference of the Independent Society of Multidisciplinary and Interdisciplinary AGI/SIGI Researchers
http://artificial-mind.blogspot.com/2012/05/news-updates-about-upcoming-first.html
http://research.twenkid.com

---

to agi

From: "Mike Tintner" <tintner@blueyonder.co.uk>
To: <agi@listbox.com>
Subject: Re: [agi] ..and Visual Individual Recognition
Date: Mon, 14 May 2012 00:18:54 +0100

I am still somewhat amazed by this kind of response - and I want to stress heavily that yours is broadly a typical AGI-er response, other than that you are going into much more detail.

You like every other AGI-er couldn't see the difference between a set of patterns and a set of patchworks. In the former, all share exactly the SAME PARTS, arranged in exactly the same structure. In the latter, all may share NO SAME PARTS, and the parts are arranged in different structures - so while each patchwork may be to the human eye somewhat similar, each of them is actually also new and different on analysis. And each patchwork may also introduce altogether new kinds of parts - a new kind of animal say. Put that another way, the patterns, you could say, are variations on exactly the same box/assortment of chocolates - cake-base chocs, say ; whereas in the patchworks, the chocolates are always radically different assortments, fruit chocs, marzipan chocs, liqueur chocs etc etc. I doubt frankly (although I haven't done research) that any non-AGI-er would have difficulty telling a patchwork from a pattern- but as I said, it ain't just you, it's every AGI-er I've discussed this with..

Somewhat similarly, you are not observing the radical differences in the features of caricatures and photos. One has a super pointy nose, the other a flat-based nose. One has majorly busy eyebrows, with the bushes focussed on different sections of the eyebrows; the other has somewhat thick brows, with the thickness differently distributed. And so on.

And the simple reality is that there aren't any programs that can see the similarity between caricatures and fotos, are there? Jeez, it's mindlbowingly difficult erough to relate photos.

And then, you come back with this 3-d projection business. A 3-d projection of a chair with a 4-strut back is NOT going to have any serious resemblance to a 3-d projection of a Chesterfield chair. 3d projections will **magnify** the differences. Abstracting the shapes won't make them any more similar either - a four-barred shape is not going to relate to a solid dimpled shape with a radically different outline; a barred shape is not going to relate to a solid shape.

What your analysis shows is that you like other AGI-ers are simply getting lost in the details and the technology.

You have to show HOW a program can tell that the features of one image can be recognized as those of another - how the caricature pointy nose say - can be recognized as a transformation of another image - the foto flat-based nose.

You never do this. You wave a massive amount of detail around, but none explains any particular transformation.

You may be right that in some cases people can't directly recognize - or mistake - a particular caricature. But it doesn't really matter. The crucial point is that, once the likeness is pointed out, they can then live v. comfortably with the caricature despite its distortions. Similarly, as I pointed out recently, people live v. comfortably with the obvious distortions of their subjects by many more or less realistic paintings . The human brain can do something that computers haven't yet been programmed to do - it can say/think: "this shape is loosely/fluidly like that one". We don't have to grasp for precision. And it's that, paradoxically at first, that is the source of our generality. Bob M

was the only one here to grasp the importance of this.

If you use geometry to define the shape of a grasping hand with precision, your robot will always be a specialist, only ever able to grasp a few special-shape objects (like current robots).

If you have a robot that can understand that a particular hand shape is to be interpreted fluidly, then it will be an AGI, able to grasp more or less any shape of object. .

VAGUE FLUID " ARTISTIC" TRANSFORMATIONS = BROAD GENERALITY/GENERAL RANGE OF FORMS

PRECISE HOMEOMORPHIC GEOMETRIC TRANSFORMATIONS -= NARROW SPECIALIST RANGE OF FORMS.

From: Todor Arnaudov <twenkid@gmail.com>
To: agi@listbox.com
Subject: Re: [agi] ..and Visual Individual Recognition
Date: Wed, 16 May 2012 02:37:11 +0200

Mike,

I'm amazed of never ending misunderstanding of all those detailed explanation, but it's a good exercise.

>If you have a robot that can understand that a particular hand shape is to be interpreted fluidly, then it will be an AGI, able to >grasp more or less any shape of object. .
>VAGUE FLUID " ARTISTIC" TRANSFORMATIONS = BROAD GENERALITY/GENERAL RANGE OF FORMS
>PRECISE HOMEOMORPHIC GEOMETRIC TRANSFORMATIONS -= NARROW SPECIALIST RANGE OF FORMS.

Just to remind you, that I've got enough of evidences for being a decent multi-talented (universal) artist, in all domains (modalities); besides a programmer with vast background in CS, engineer with strong mathematics background; been a journalist, also a philosopher and has a solid background in many other fields. Sorry to self-promote, but really I don't know of a more-versatile ("general") intelligence person.

Good for you to believe you understand art and images better and I'm yet another "AGI-er" with so and so confusions. You're right most AGI-ers are not artists, and don't think "fluidly" in terms of art, but some are and see both or many worlds.

>And the simple reality is that there aren't any programs that can see the similarity between caricatures and fotos, are there?

As I've said - that's a shame for Computer Science, AI and AGI. Also for me - because I haven't even started to write my thinking machines yet, I can't afford to sit and focus, still on the way of funding myself while the prototype won't be complete enough to be commercial...

>Jeez, it's mindlbowingly difficult erough to relate photos.

It's not.

>And then, you come back with this 3-d projection business. A 3-d projection of a chair with a 4-strut back is NOT going to have
>any serious resemblance to a 3-d projection of a Chesterfield chair. 3d projections will **magnify** the differences. Abstracting
>the shapes won't make them any more similar either - a four-barred shape is not going to relate to a solid dimpled shape with a
>radically different outline; a barred shape is not going to relate to a solid shape.

Mike, first of all it's not 3D projection, but a 2D-projection of 3D world (vision with a single camera or a single eye), and 3D-reconstruction from a 2D image or video, or/and a pair of 2D images/videos, produced by 2D projection of a 3D input.

The reconstructed 3D-model at different resolutions and different focus is the thing to compare, the model, not the projection, rotation etc.

If you recall the generalization example about the chairs, what's general and essential in all chairs is **the plane which is parallel to the ground, and has a supporting vector which is**

perpendicular to the ground.

Every input/entity which has those simple features in its 3D-reconstructed model will look similar in this regard to all other entities which have that feature, no matter how many legs it has, does it have a back or not, is it 1m or 100 m tall, does it have curves or not - all those don't matter.

It doesn't matter how you call it, everything having those will be similar to each other, just because it's always found there, on each analysis of the input a plane and a support will be found, while the other features would be different.

Yes, the details are different, if the 3D-object is rotated, the 2D-projection will differ a lot from a frame to frame, but the 3D-model will be the same, rotated or translated 3D-model is the same to its basis vector, that's how it's compared.

A "barred shape", "solid shape" or "textured shape", they all have a level of representation when their texture doesn't matter.

Again, something about generalization that you confuse:

--> Of course if a similarity is in demand to be found in N objects, and they all have different texture, color, appearance etc., then *that", what is *different*, apparently is not what's "general", what groups those object together.

--> Something else should be the one, and it often is very simple - otherwise it won't fit in all cases.

That "else" is another representation, another resolution, another feature. E.g. a voxel representation. Or a "bounding box". **All the rest might be different, but if at a level of "bounding box" two objects may match, then they are "similar", or at that level of abstraction - "the same".**

-- At a certain level of compression/generalization, everything is the same, one of the core problems in AGI is the engine that plays with forming, selection and focus on those levels.

Overall, to recognize anything as "similar", on the road there's always something that matches completely. At the final step of caricature recognition that is the person, the original photo etc.

>What your analysis shows is that you like other AGI-ers are simply getting lost in the details and the technology.
>You have to show HOW a program can tell that the features of one image can be recognized as those of another - how the >caricature pointy nose say - can be recognized as a transformation of another image - the foto flat-based nose.
>You never do this. You wave a massive amount of detail around, but none explains any particular transformation.

I do it, you miss it. Voxels and pixels. Resolution reduction - that's all from scaling, resizing, reducing the number of pixels, contrast etc. Check out some Computer vision and image processing.

What's missing is one thing: INTEGRATION.

> You wave a massive amount of detail around, but none explains any particular transformation.

Mike, I wave that "massive" amount of detail, because I see it, and explained you the transformation, what's simpler than increase and decrease. And this is not a massive amount of detail, just a few bytes.

If I was to explain it in code, you would have said the same because you're not a programmer, and if I had the code, I'd not explain it here. Besides, this comparison and recognition requires not only code in C++, the samples are the "code" (as Boris has mentioned also), the execution of the actual comparisons would be verbose; the directions and principles are obvious, though.

>But it doesn't really matter. The crucial point is that, once the likeness is pointed out, they can then live v. comfortably with the
>caricature despite its distortions. Similarly, as I pointed out recently, people live v. comfortably with the obvious distortions of
>their subjects by many more or less realistic paintings . The human brain can do something that computers haven't yet been
>programmed to do - it can say/think: "this shape is loosely/fluidly like that one". We don't have to grasp for precision. And it's
>that, paradoxically at first, that is the source of our generality. Bob M was the only one here to grasp the importance of this.

Namely that "loosely/fluidly" is **that loss of resolution and details which allow for us to accept those distorted pictures as references to the others.**

That "fluidity" you talk about is everything that lays between the boundaries defined by the "patterns", it doesn't need to be explicitly defined - everything that's not forbidden explicitly is allowed. The "patterns" are limits, and they are not static. Those boundaries are expressed in those details I mentioned, which are just a few bits.

A raw image may cost 10000 KB in computer memory (not very high resolution scan), **the essential features (aside the samples to compare them with) may fit in 100 bytes. Just resize the picture, reduce it to BW, or better - ask an artist to trace it and record it as a vector image - a few bytes. All the rest is "decoration".**

Those bytes are encoded in the details that I mentioned.

>If you use geometry to define the shape of a grasping hand with precision, your robot will always be a specialist, only ever able >to grasp a few special-shape objects (like current robots).

Well, even trivial higher differential geometry includes "curves", they are "fluid". Have you heard of "parametric curves", such as B-splines and Bezier curves. Any "fluidity" can be defined in complex enough curve, having enough of parameters and ranges.

Cheers,
Todor "Tosh" Arnaudov
http://research.twenkid.com
http://artificial-mind.blogspot.com

http://artificial-mind.blogspot.bg/2012/01/colour-optical-illusions-are-effect-of.html

Sunday, January 1, 2012

Colour Optical Illusions are the Effect of the 3D-Reconstruction and Compensation of the Light Source Coordinates and Light Intensity in an Assumed 2D Projection of a 3D Scene

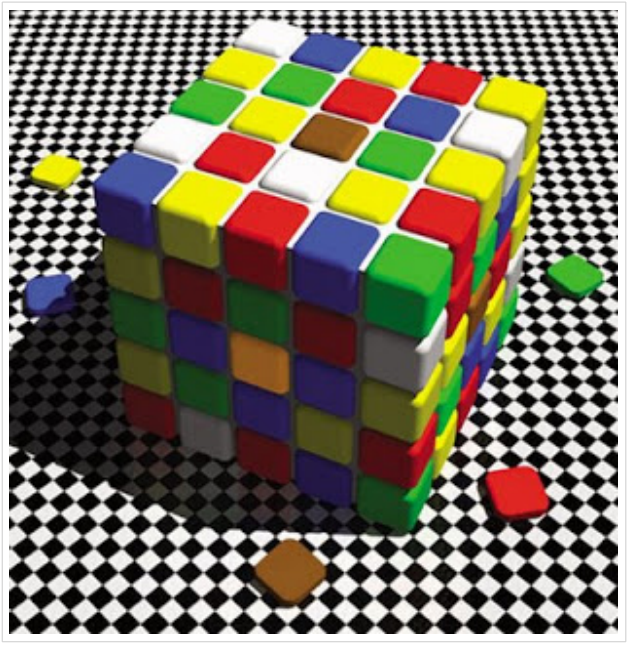

1) In the 2D image the pixels of the brown square in the middle on the top side of the cube is actually the same pixel color as the orange square in the side which looks in the shadow part of the cube. Notice the picture below, there's more contrast and the square looks darker.

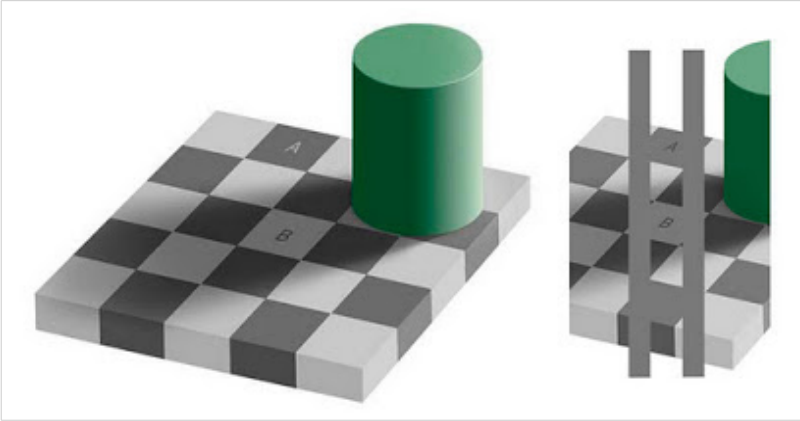

2.) Squares A and B in the 2D picture have the same absolute shade of gray.

Looks mysterious, there are claims that these are "prewired" bugs. My explanation is pretty straight, and these color illusions are rather a feature than a bug. It seems as a bug only in perverse cases as in the illusions, and only before the reason is understood:

- ## Introduction
- **Light source coordinates, lights colors and reflections are ambiguous**

  2D images, and even the "stereo" images from our eyes from the real world are ambiguous regarding the light source and the colors, because the scenes are usually partially observable and because there's only perception of absolute color - "RGB" and luminance (rods and cones). Besides light also has color and there can be reflections and shadows which can change both luminance and color of the affected areas. Another ambiguous components are the texture, reflectiveness and transparency of the objects, which also alters the low level RGB properties.

  "Light sources" are not perceived at the low level input - it is a construct created in order to explain the patterns of differences in luminance and color which are observed. The places with high luminance are seen as "lit", the ones with lower - "shadowed", and the gradients of luminance seem as light following shapes, dispersing, reflecting etc. What is actually "following the shapes of the objects", dispersing, reflecting etc. is the luminance of the "pixels" in the 2D image.

  It's similar with transparency, reflectiveness etc. - these are stable correlation observed in many samples, which gave birth of the prediction that given such correlation, the object/that part of the 2D image or reconstructed 3D scene should have a property "transparency" etc.
-
- **3D reconstruction from a 2D image is ambiguous**

  Cognitive hierarchy gets samples of 2D images which are projections of 3D objects. Even more - stereo samples allow it to do its kind of "triangulation" and reconstruct distances so that we perceive depth if we look with both eyes.

  It goes even further - if one has grown with two healthy eyes, distance can be measured even if you close one of them. And even more: humans see 3D space in 2D projections on a flat surface (a picture or electronic screen) using only one eye!

  That's possible because of the "aggressively" repetitive perception of correlations between the 2D projections on the retina and the perspective changes of the image after controlled and predictable changing the viewing angle or/and object coordinates.

  The correlations of the change of luminance and color depending on the light source (changes in contrast, luminance and color) are also observed.

  A lot of 2D clues for the 3-rd dimension of the objects and the position and intensity of light are extracted and "fixed" so that future 2D images are perceived as 3D even if having just connected parallel lines and other minute details characteristic for perspective transformation.

I've discussed a little bit on this back in 2004 in my [Teenage] Theory of Universe and Mind, Part 4:
The section starts from the last line at p.30 in this translation.
// The picture there seems not to appear in the document, it's the one below: //

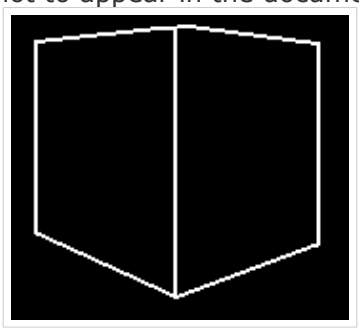

Mind sees 3D space and light source(s) everywhere

The price of the simplicity of 3D-clues and light is that these repetitive correlations cause mind to expect 3D space and light source(s) even in the 2D flat images. Artists use to draw 2D projections of a 3D space, both photographs and video do etc, the clues are constantly reinforced.

It seems that the highest consciousness level can hardly inhibit 3D perception/integration on demand and see 2D images as 2D if any 3D-clues intervene and confuse it.

[Edit/note from: 23/7/2012 - Therefore 3D is more powerful, it has more impact to mind and seems to come first in such perceptions, and that's reasonable - because humans initially sense the world as 3D with two eyes, and all our physical muscular reactions, which in essence are just sequences of coordinate adjustments, are in 3D, not in 2D, the real world is perceived as not flat. Follow the blog for continuation on the topic.]

So...

Why the brown square looks orange?

Because the 3D-reconstructed scene seems to display a cube which is lit from the top and the side which is directed towards the screen is in the shadow, there are too much clues suggesting this, so mind believes.

However if the square in the shadow has the same absolute 2D retinal color ("RGB") as the lit square, and it is still in a shadow, then the original unlit color of the shadowed square should have been **orange**, because orange at this relatively lower luminance turns into the brown - the same like the top square was.

The case with the gray squares and the cylinder on the other picture is similar. The connections and borders added on the right illustration destroy the clues suggesting that these particular squares are part of the 3D scene with the cylinder and the light source behind it, because they don't follow the pattern of light. There are not clues that they are behind the scene either (being partially covered - there are no interrupted lines), so the "Z-buffer" puts them in front of the other elements. Now there is not a light source and a shadow over those squares, and their absolute retinal RGB values are not adjusted differently.

Mystery unveiled... :)

Hypothesis

- A quick hypothesis might be that young infants that don't see depth yet may not have this optical illusion.
- However I'm not sure about that - I suspect that babies should understand that light sources change color and luminance of the areas they impact long before they have robust 3D-perception, 3D is not needed, so I guess they may have clues of light and shadow interaction and see the optical illusion even before they see 3D.

It will be just a 2D optical illusion, because I suspect brain probably does a partial 2D light-color reconstruction before it goes into 3D. (It would be cool if I could conduct some kind of experiments, but not possible yet.)

-

2D-projections of 3D scenes are "2D 1/2" or "Quasi 3D" anyway, and that indeed is the reason for the impossible perspective drawings of Escher and other artists.

3D-reconstruction is ambiguous and these drawings suggest that it's done locally, based on local 2D-clues. Globally the picture might be absurd, but mind fails in focusing in on all 3D-clues, so it sees correct perspective in the different elements.

Why can't we correct this illusion consciously?

One reason I'd say is that is not really an illusion.

Another answer is in the recent series I posted about higher-lower level feedback/reflection in the cognitive hierarchies which I think is much weaker/narrower than the feedforward.

Higher level understands some of the "mistakes" and "delusions" (wrong predictions) of the lowers if takes details separately out of their context, but it cannot correct them if they are too many levels below and if there are too many *correct* predictions at the lower levels. The clues suggesting there's no illusion and the predictions *work* are reinforced for too many cases to be dismissed and altered.

In the examples above everything in the low level input suggests that there is no illusion: all other correlations between absolute pixel values confirm that these are pictures of 3D-scenes with their light sources and shadows.

Another explanation might be that high levels lack direct access to the lowest levels - retinal RGB-luminance values.

...

To be continued...

* Thanks to a show on Discovery Channel about optical illusions that made me think about this yesterday and realize this mechanism.

** Happy New Year, wishing 2012 to be more productive in AGI/SIGI.

- Оптични илюзии с цветове - възстановяване на триизмерните координати на източника на светлина в двуизмерна проекция на тримерна сцена.

Posted by Todor "Tosh" Arnaudov at 8:01 AM