

Toward Evolutionary and Developmental Intelligence

Kenji Doya^a & Tadahiro Taniguchi^b

a: doya@oist.jp

Neural Computation Unit, Okinawa Institute of Science and Technology Graduate
University

1919-1 Tancha, Onna, Okinawa, 904-0495, Japan

b: taniguchi@ci.ritsumeikan.ac.jp

College of Information Science and Engineering, Ritsumeikan University

1-1-1 Noji Higashi, Kusatsu, Shiga, 525-8577, Japan

Abstract:

Given the phenomenal advances in artificial intelligence in specific domains like visual object recognition and game playing by deep learning, expectations are rising for building artificial general intelligence (AGI) that can flexibly find solutions in unknown task domains. One approach to AGI is to set up a variety of tasks and design AI agents that perform well in many of them, including those the agent faces for the first time. One caveat for such an approach is that the best performing agent may be just a collection of domain-specific AI agents switched for a given domain. Here we propose an alternative approach of focusing on the process of acquisition of intelligence through active interactions in an environment. We call this approach *evolutionary and developmental intelligence* (EDI). We first review the current status of artificial intelligence, brain-inspired computing and developmental robotics and define the conceptual framework of EDI. We then explore how we can integrate advances in neuroscience, machine learning, and robotics to construct EDI systems and how building such systems can help us understand animal and human intelligence.

Introduction

Spurred by successful scaling of deep learning (DL) [1,2] to huge complex data sets, artificial intelligence (AI) today has achieved supra-human performance in specific domains like visual object recognition [3] and game playing [4-6]. The next focus in AI

research is to flexibly find solutions to novel tasks, under the concept of artificial general intelligence (AGI) [7]. There has been much discussions on how to define, design and evaluate AGI agents (e.g., <http://cadia.ru.is/workshops/aegap2018/>). One approach is to set up a variety of tasks, including those the agent faces for the first time, and test how AI agents perform (e.g., <https://www.general-ai-challenge.org>). One caveat for such an approach is that the best performing agent may be just a collection of domain-specific AI agents switched for a given domain.

A critical problem in such approaches to AGI is their focus on the achieved performance after learning. We advocate an alternative approach to focus on the process of acquisition of intelligence. We call this approach *evolutionary and developmental intelligence* (EDI). Animals, especially humans, can learn relevant features in the sensorimotor signals in an unsupervised way, build internal models of the world including the agent itself, find a variety of action policies, and further set up new goals of actions. Such capabilities for incremental learning were not given by an external designer but established on their own through evolutionary search for the fitness in the environment. In other words, we should pay more attention to the process of acquisition of intelligence through development of individual AI agents and evolution of AI architectures and algorithms. Here we consider what are missing in the current implementation of AI agents and how we can evolve AI agents that develop like animals and humans, by exploiting and extending our knowledge from neuroscience, machine learning and robotics.

Why evolutionary and developmental approach?

Despite the impressive success stories by DL, there are still major gaps between what machine learning today can offer and what humans, even children can do [8,9]. Most notably, data efficiency and energy efficiency. The former is based on our capability of inference by analogy and compositional use of knowledge and skills. An even more fundamental difference is whether an agent is designed or instructed to perform a certain task, or can find its own goals or problems. These gaps between today's AI and human cognition urges us to search for clues and principles in the brain [10].

Autonomy and evolvability

Current approach to AI is for a human developer to define the problem to be solved, collect relevant data, design a neural network architecture or a probabilistic graphical model, and then apply a learning algorithm for solution. Here we advocate a totally

different approach for creating autonomous intelligent agents. Before asking an agent to do something particular, let agents acquire the capability of survival and reproduction, which are the fundamental features for living and evolvable agents [11,12] (**Box 1**). In physical robots, that requires basic sensory-motor mechanisms for capturing energy sources, avoiding dangers, and performing reproduction, guided by innate behaviors and learning by primary rewards. In software agents, survival means continuing to be utilized and reproduction means proliferation of the copies. On top of such autonomy, each agent explores the environment to incrementally acquire wider varieties of sensory-motor features and build dynamic models of the world including its peers and itself. This process is guided by learning with intrinsic rewards [13-15]. If such an agent is to perform a certain task which a human desires, it is guided by an additional social rewards, as we would for training animals or educating children. This is certainly a long way around for solving a well-defined task and may appear like a daydream. We argue, however, that this is feasible and the most certain way for building autonomous agents with human-level flexibility. Setting particular goals on top of the basic principle of survival and reproduction may also help avoiding the headache of programming common senses, such as do not destroy oneself or do not do things hated by others. It may take millions or billions of years if we follow the way humans evolved, but there are many shortcuts and accelerations we can make by utilizing the knowledge of neuroscience and the advances in information technology.

Box 1: Evolution of rewards and polymorphisms in artificial agents.

The aim of experiments with *Cyber Rodent* robots [11] (Figure 1A) was to test whether a colony of robots with the capability of battery recharging and software exchange by infrared (IR) communication could acquire their own reward functions for the sake of survival and reproduction. Each robot had vision and proximity sensors and two wheels for navigation, and reinforcement learning controllers for foraging and mating, which were switched by a top-level neural network with sensory inputs including internal battery level. They exchanged their *genes* (weights of the top-level network and reward function networks, and reinforcement learning parameters) through IR communication. The probability of selection in the next generation was proportional to the parent's battery level and mutation by small random noise was applied. Over 40 generations of evolution in simulation, distinct reward functions for the sight of a battery pack and another robot were obtained [12]. In some of the colonies, individuals with

distinct mating strategies co-existed; *foragers* who mate only after fully charged and *trackers* who opt for frequent mating even when the battery level is low. Further analyses showed that these subtypes had distinct genotypes and were evolutionarily stable [16].

Figure 1: (A) *Cyber rodent* robot colony. (B) The evolved reward functions for the vision of a battery pack (left) and another robot’s face (right). (C) Sub-populations of robots taking *forager* and *tracker* strategies (left) with distinct genotypes (right).

Developmental Psychology:

One of the remarkable findings from the human genome project is that the number of genes in humans is about 30 thousand, which is much fewer than the number of neurons or synapses in the brain. This means that most of the information stored in the brain is acquired from the environment, while genes provide efficient mechanisms for acquiring information. Sensory-motor interaction with the environment is a critical requirement of human cognitive development. While there appear to be innate mechanisms for basic cognition, such as recognizing facial expressions [17], most of the knowledge and skills are acquired by sensory-motor interaction with the physical and social environment. Infants as young as 2 months old can detect unusual physical contingencies [18] and 6 months old can discriminate the intentions of animated agents [19]. Such capabilities, termed intuitive physics and intuitive psychology, are the basis for our everyday thinking and behaviors, and therefore indispensable for artificial agents working in the human society [9].

Evolutionary and developmental robotics

Originating from Piaget’s concept of *constructionism* [20], the major focus of evolutionary and developmental robotics has been how an embodied agent can acquire sensory perception, motor control, and higher cognitive capabilities through bottom-up unsupervised interactions with the physical world, including other robots and humans [21,22].

By incorporating advances in probabilistic models and deep learning, there have been much progress in developmental robotics. For example, Taniguchi and colleagues developed SpCoSLAM in which a mobile robotic incrementally acquires multi-modal probabilistic models of visual objects, spoken words, and their locations for navigation [23]. Tani demonstrated that cognitive functions like sequences of

motion primitives and compositionality of words can emerge through embodied interactions using deep neural networks implementing the principle of predictive coding [24,25]. A major feature of these approaches is that symbol-like representations emerge through sensory-motor interactions with the world [26,27], which is opposite to the situation of “symbol grounding” that hampered classic symbolic AI [28].

Recent advances and the way forward

Now we outline how such evolutionary and developmental AI systems can be practically constructed by building on and further advancing neuroscience, machine learning and robotics.

Neuroscience:

The capability of learning is a product of evolution. While single-cell organisms or tiny worms have varieties of mechanisms for learning and memory, the mammalian brains have acquired distinct mechanisms for learning; error-driven supervised learning in the cerebellum, reward-guided reinforcement learning in the amygdala and basal ganglia, episodic memory in the hippocampus, and Bayesian inference and representation learning in the cerebral cortex [29-31].

A critical component of human intelligence is to learn internal models of the world and to run simulations of the world for estimating the causes of sensory perception, planning actions to achieve desired goals, and running thought experiments of arbitrary situations. The neural mechanisms of such *mental simulation*, or model-based inference and control, is now being revealed using advanced imaging and computational analyses [32-37] (Box 2). While computers are very good at running simulations and searches, how to build and select models of appropriate levels of abstraction and concreteness, and how to direct searches to the right width and depth are still open problems. Understanding of the neural substrates of mental simulation at the whole brain and local circuit levels would provide vital insights for the design of human-like flexible intelligent agents.

Box 2: Neural substrates of mental simulation.

Mental simulation, which we define as the brain’s process using action-dependent state transition models, is a critical component of intelligence. Given recent advances in neural imaging, experimental interrogation of the neural circuits that realize mental simulation is now becoming feasible. In the functional

MRI experiment using the ‘grid sailing task,’ subjects planned ahead zig-zag paths to the goal location using pre-trained key maps, i.e. key-press dependent cursor transition models [35] (Figure 2A). The brain activity during the pre-movement delay period suggested the involvement of a global network linking the parietal, premotor and prefrontal cortices, which can provide spatial and motor working memory, with the cerebellum and basal ganglia, which can provide forward models and value functions, in mental simulation. For finer analysis of the neural circuit of model-based inference, Funamizu et al. performed two-photon imaging of the parietal cortex while mice performed navigation under uncertain sensory feedback [36] (Figure 2B). Blockade of parietal cortex impaired estimation of the goal position under missing auditory feedback. Decoding of the population codes of parietal neurons showed that the representation of goal position was updated even without auditory feedback by action-dependent predictive models.

Figure 2: (A) The grid sailing task [35] in which a subject tries to move the cursor, that can move only in three directions, from the start to the goal on a grid. Subjects’ performance improved with pre-learned action-dependent state transition model and pre-start delay time, which is the behavioral evidence of mental simulation. (B) Enhanced activities in the cortical, cerebellar, and basal ganglia areas were observed during the pre-start delay time. (C) Two-photon imaging of parietal cortical neurons of mice in an auditory virtual environment [36]. (D) Decoding by probabilistic population codes revealed that the representation of the goal distance were updated even without sensory feedback in consistence with the mice’s own locomotion.

Machine learning:

While early successes of deep learning (DL) was by supervised learning using labelled training data [38], recent developments in DL focus on unsupervised or self-supervised learning of deep generative models, such as variational autoencoders (VAE) [39] and generative adversarial networks (GAN) [40]. By incorporating recurrent connections, such as the long short-term memory (LSTM) [41], such deep generative models can also predict and generate spatio-temporal dynamics, such as speech, language and movements. One domain of active research is meta-learning for automatically selecting network architectures and parameters [42-45]. In

reinforcement learning, there are demonstrations that by training a single network with multiple tasks, the latent structures relevant for achieving the tasks can be captured in the hidden units through bottom-up interactions with the environment [46,47].

Probabilistic programming languages [48], which allow flexible designs of probabilistic models and derivation of their inference algorithms, are now incorporating deep neural networks as generative models [49], which is a favorable function for acquisition of world models through real sensory-motor interactions. For constructing multi-modal generative models in a modular and flexible way, Nakamura et al. proposed SERKET, a framework for connecting probabilistic generative models by efficient inter-module communication [50].

Robotics and human-robot interaction

For robotic agents to develop internal models of the physical world and human behaviors, sensorimotor interaction with its environment and social interactions with humans are essential. Fluid use of language, for example, requires not only lexical grammatical knowledge but also understanding of the physical context, such as what the speaker is doing now, and inference of the speaker's intention. When a robot tries to understand a command like "go to the kitchen, and take me a bottle of water," the robot has to deal with object and place concepts, action, syntax, planning and so on. This means that an actual language learning involves multimodal concept learning, action learning, syntax learning and so on.

A critical issue in doing all these through simple sensorimotor interactions is the time needed for learning, especially with the notorious data-hungriness of deep learning. However, robots and computers have their specific advantages of replaceability and network communication. While human bodies differ a lot across individuals, robots can be manufactured physically similar so that it is practical to collect sensorimotor data from multiple copies of robots to accumulate large amount of data for learning. In other words, telecommunication across brains like telepathy and copying the learned neural network like brain transplant, which are technically and ethically difficult in humans, are quite easy in robots. Smartphones can be the media for collecting data of visual, auditory, and linguistic interactions in everyday human life, either by giving them minimal actuators [51] or by using their owners as mobile caretakers.

Conclusion

An amazing fact about the brain and human cognition is that specialized neural circuits as well as the mechanisms for integrating those networks could be realized through evolutionary search, by exploiting any usable features of biophysics of neurons and statistical dynamics of networks. There is no known engineering solution to building such complex heterogeneous systems other than evolutionary optimization. For example, in protein engineering, *directed evolution* has seen successes in finding complex molecules having desired functions [52]. Designing AI systems to be evolvable, rather than hard-coded by human intuition and theorization, can be a rational practical choice.

In order not to repeat the whole history of life, it is possible and practical to set the starting points of evolutionary search to what are already known to work. Unlike real lifeforms on earth, artificial agents can perform Lamarckian way of copying learned behaviors through the internet with thousands of peers anywhere. In the field of visual object recognition, most researchers had believed that the use of human-engineered features are the best way, until end-to-end data-driven learning by deep neural networks outperformed them [3]. In building AGI, although it might sound daydream or waste of time, mimicking the evolutionary and developmental paths of human cognition may be a practical solution, given those possible shortcuts and accelerations.

Along the way of such an endeavor, we should encounter many unexpected abnormal or suboptimal performances of evolving/developing agents. Such examples, however, could be helpful models to understand the mechanisms of genetic or developmental cognitive disorders. The forms of intelligence that are found by evolution may be different from those of humans, like those of birds or octopi, or nothing like those on earth, depending on the given constraints. Comparing the performances of artificial agents with humans, especially in human-robot interactions, can be helpful tools to clarify what are missing in our understanding of human behavior and cognition.

Acknowledgements

This work was supported by Ministry of Education, Culture, Sports, Science and Technology KAKENHI Grants 23120007, 16K21738, 16H06561 and 16H06563, and research support of Okinawa Institute of Science and Technology Graduate University to KD, and KAKENHI Grant 16H06569 and JST CREST Grant JPMJCR15E3 to TT.


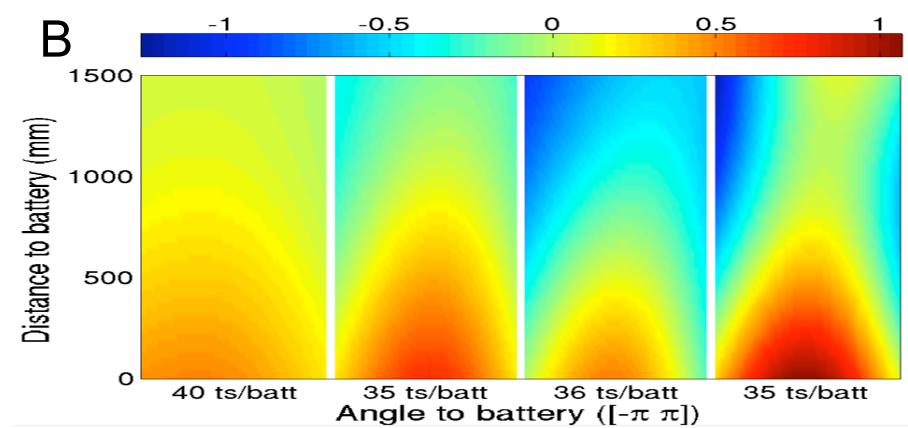

References

1. LeCun Y, Bengio Y, Hinton G: **Deep learning**. *Nature* 2015, **521**:436-444.
2. Schmidhuber J: **Deep learning in neural networks: an overview**. *Neural Netw* 2015, **61**:85-117.
3. Krizhevsky A, Sutskever I, Hinton GE: **ImageNet classification with deep convolutional neural networks**. In *Advances in Neural Information Processing Systems*. Edited by; 2012:1090–1098. vol 25.]
4. Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, et al.: **Human-level control through deep reinforcement learning**. *Nature* 2015, **518**:529-533.
5. Silver D, Huang A, Maddison CJ, Guez A, Sifre L, van den Driessche G, Schrittwieser J, Antonoglou I, Panneershelvam V, Lanctot M, et al.: **Mastering the game of Go with deep neural networks and tree search**. *Nature* 2016, **529**:484-489.
6. Silver D, Schrittwieser J, Simonyan K, Antonoglou I, Huang A, Guez A, Hubert T, Baker L, Lai M, Bolton A, et al.: **Mastering the game of Go without human knowledge**. *Nature* 2017, **550**:354-359.
7. Goertzel B, Pennachin C (Ed): *Artificial General Intelligence*: Springer; 2006.
8. Lake BM, Salakhutdinov R, Tenenbaum JB: **Human-level concept learning through probabilistic program induction**. *Science* 2015, **350**:1332-1338.
9. Lake BM, Ullman TD, Tenenbaum JB, Gershman SJ: **Building machines that learn and think like people**. *Behav Brain Sci* 2017, **40**:e253.
10. Hassabis D, Kumaran D, Summerfield C, Botvinick M: **Neuroscience-inspired artificial intelligence**. *Neuron* 2017, **95**:245-258.
11. Doya K, Uchibe E: **The Cyber Rodent Project: Exploration of adaptive mechanisms for self-preservation and self-reproduction**. *Adaptive Behavior* 2005, **13**:149-160.
12. Elfwing S, Uchibe E, Doya K, Christensen HI: **Darwinian embodied evolution of the learning ability for survival**. *Adaptive Behavior* 2011, **19**:101-120.
13. Kaplan F, Oudeyer P-Y: **In search of the neural circuits of intrinsic motivation**. *Frontiers in Neuroscience* 2007, **1**:225-236.
14. Uchibe E, Doya K: **Finding intrinsic rewards by embodied evolution and constrained reinforcement learning**. *Neural Networks* 2008, **21**:1447-1455.
15. Baldassarre G, Stafford T, Mirolli M, Redgrave P, Ryan RM, Barto A: **Intrinsic motivations and open-ended development in animals, humans, and robots: an overview**. *FiCS* 2014.
16. Elfwing S, Doya K: **Emergence of Polymorphic Mating Strategies in Robot Colonies**. *Plos One* 2014, **9**.
17. Demirir Y, Meltzoff A: **The robot in the crib: A developmental analysis of imitation skills in infants and robots**. *Infant Child Dev* 2008, **17**:43-53.
18. Hespos SJ, Baillargeon R: **Young infants' actions reveal their developing knowledge of support variables: Converging evidence for violation-of-expectation findings**. *Cognition* 2008, **107**:304-316.
19. Hamlin JK, Wynn K, Bloom P: **Social evaluation by preverbal infants**. *Nature* 2007, **450**:557-559.
20. Flavell JH: *The Developmental Psychology of Jean Piaget*: Van Nostrand Reinhold Company; 1963.
21. Pfeifer R, Scheier C: *Understanding Intelligence*: MIT Press; 2001.
22. Cangelosi A, Schlesinger M: *Developmental robotics: From babies to robots*: MIT Press; 2015.
23. Taniguchi A, Hagiwara Y, Taniguchi T, Inamura T: **Online spatial concept and**


- lexical acquisition with simultaneous localization and mapping.** In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Edited by; 2017:811-818.
24. Tani J: *Exploring robotic minds: actions, symbols, and consciousness as self-organizing dynamic phenomena*: Oxford University Press; 2016.
 25. Tani J, Ito M, Sugita Y: **Self-organization of distributedly represented multiple behavior schemata in a mirror system: reviews of robot experiments using RNNPB.** *Neural Networks* 2004, **17**:1273-1289.
 26. Taniguchi T, Nagai T, Nakamura T, Iwahashi N, Ogata T, Asoh H: **Symbol emergence in robotics: a survey.** *Advanced Robotics* 2016, **30**:706-728.
 27. Taniguchi T, Ugur E, Hoffmann M, Jamone L, Nagai T, Rosman B, Matsuka T, Iwahashi N, Oztop E, Piater J, et al.: **Symbol emergence in cognitive developmental systems: A survey.** *IEEE Transactions on Cognitive and Developmental Systems* 2018:1-1.
 28. Harnad S: **The symbol grounding problem.** *Physica D: Nonlinear Phenomena* 1990, **42**:335-346.
 29. Doya K: **What are the computations of the cerebellum, the basal ganglia, and the cerebral cortex.** *Neural Networks* 1999, **12**:961-974.
 30. Doya K: **Complementary roles of basal ganglia and cerebellum in learning and motor control.** *Current Opinion in Neurobiology* 2000, **10**:732-739.
 31. Doya K: **Bayesian brain: Bayesian inference for functional modeling and data analysis of the brain.** In *Institute of Statistics and Mathematics November 8, 2006; Tokyo*: 2006.
 32. Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ: **Model-based influences on humans' choices and striatal prediction errors.** *Neuron* 2011, **69**:1204-1215.
 33. Suzuki S, Harasawa N, Ueno K, Gardner JL, Ichinohe N, Haruno M, Cheng K, Nakahara H: **Learning to simulate others' decisions.** *Neuron* 2012, **74**:1125-1137.
 34. Donoso M, Collins AG, Koechlin E: **Foundations of human reasoning in the prefrontal cortex.** *Science* 2014, **344**:1481-1486.
 35. Fermin AS, Yoshida T, Yoshimoto J, Ito M, Tanaka SC, Doya K: **Model-based action planning involves cortico-cerebellar and basal ganglia networks.** *Sci Rep* 2016, **6**:31378.
 36. Funamizu A, Kuhn B, Doya K: **Neural substrate of dynamic Bayesian inference in the cerebral cortex.** *Nat Neurosci* 2016, **19**:1682-1689.
 37. Doll BB, Bath KG, Daw ND, Frank MJ: **Variability in dopamine genes dissociates model-based and model-free reinforcement learning.** *J Neurosci* 2016, **36**:1211-1222.
 38. Salakhutdinov R, Hinton G: **An efficient learning procedure for deep Boltzmann machines.** *Neural Comput* 2012, **24**:1967-2006.
 39. Kingma DP, Welling M: **Auto-encoding variational Bayes.** In *International Conference on Learning Representations (ICLR)*. Edited by; 2014.
 40. Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y: **Generative adversarial nets.** In *Neural Information Processing Systems (NIPS)*. Edited by; 2014.
 41. Hochreiter S, Schmidhuber J: **Long short-term memory.** *Neural Computation* 1997, **9**:1735-1780.
 42. Thrun S, Pratt L (Ed): *Learning to Learn*: Springer; 1998.
 43. Doya K: **Metalearning and neuromodulation.** *Neural Networks* 2002, **15**:495-506.
 44. Schweighofer N, Doya K: **Meta-learning of reinforcement learning.** *Neural Networks* 2003, **16**:5-9.
 45. Real E, Moore S, Selle A, Saxena S, Suematsu YL, Le Q, Kurakin A: **Large-scale evolution of image classifiers.** In *International Conference on Machine Learning (ICML)*. Edited by; 2017.

46. Song HF, Yang GR, Wang XJ: **Reward-based training of recurrent neural networks for cognitive and value-based tasks.** *Elife* 2017, **6**.
47. Wang JX, Kurth-Nelson Z, Kumaran D, Tirumala D, Soyer H, Leibo JZ, Hassabis D, Botvinick M: **Prefrontal cortex as a meta-reinforcement learning system.** *Nat Neurosci* 2018.
48. Ghahramani Z: **Probabilistic machine learning and artificial intelligence.** *Nature* 2015, **521**:452-459.
49. Tran D, Hoffman MD, Sauros RA, Brevdo E, Murphy K, Blei DM: **Deep probabilistic programming.** In *International Conference on Learning Representations (ICLR)*. Edited by; 2017.
50. Nakamura T, Nagai T, Taniguchi T: **SERKET: An architecture for connecting stochastic models to realize a large-scale cognitive model.** *Front Neurorobot* 2018, **12**:25.
51. Wang J, Uchibe E, Doya K: **Adaptive baseline enhances EM-based policy search: Validation in a view-based positioning task of a smartphone balancer.** *Front Neurorobot* 2017, **11**:1.
52. Gibney E, Van Noorden R, Ledford H, Castelvechi D, Warren M: **'Test-tube' evolution wins Chemistry Nobel Prize.** *Nature* 2018, **562**:176.

Figure 1



C



D




Figure 2

