美国产业动态

2019 年 4 月 4 日　　第 8 期

中国电子信息产业发展研究院（赛迪研究院）

人工智能加速逼近临界点

2019 年 1 月 30 日，国际著名市场研究咨询机构 CB 视点公司（CB Insights）发布《2019 年人工智能发展趋势》，报告从行业采纳度（包括初创企业发展势头、媒体关注度、客户采纳度）和市场实力（包括市场规模预测、融资能力、经济效益）两个维度，将 25 项人工智能应用划分为必要性、

暂时性、威胁性、试验性[1]四个发展趋势。报告认为，目前人工智能正在各个行业进行商业化布局，人工智能产业深刻影响了各个产业的发展走向。

一、“威胁追踪”——基于人工智能的一种安全防御方法，正日益受到重视。所谓“威胁追踪”，是指主动发现恶意活动，而不仅仅是对警报或入侵事件被动做出反应。据全球公共数据泄露数据库统计，2017 年全球共有 26 亿份数据被泄露，2018 年上半年则多达 45 亿份。数据泄露激增的背后是计算能力和算法的进步，让黑客攻击变成现实。仅仅对网络攻击做出反应是不够的，主动利用机器学习“追踪”黑客的威胁，受到大型公司的重视。例如，亚马逊收购威胁追踪初创公司 Sqrrl，用于搜索黑客；2018 年黑莓收购专注于威胁追踪的人工智能初创企业 Cylance。

二、开源框架大幅降低人工智能技术的门槛，初创小企业大量参与。低代码量和无代码开源软件具有低成本、低要求、高灵活的特点，迅速降低了人工智能的运用开发门槛，为初创企业提供了平等的技术平台，可与科技巨头同台竞技。如谷歌的 TensorFlow 机器学习库、蒙特利尔学习算法研究所的 Theano 开源库以及脸书的 PyTorch 开源机器学习平台。

三、“端智能”利用人工智能算法可实现无网络条件下正常信息处理及决策。人工智能算法在智能手机、汽车、可

[1] 必要性：人工智能技术在市场已得到广泛应用、认可；暂时性：只在某些领域得到应用，但前景不明确；威胁性：将取代部分传统制造业；试验性：尚处在试验阶段。

穿戴设备等终端设备的应用，可以让终端设备在无网络条件下进行正常的信息处理，并对情况做出更快的反应。例如，在没有互联网连接的情况下，自动驾驶汽车也可以对路况做出实时响应，降低因决策时间延迟造成的生命危险。目前，英伟达、高通和苹果以及许多新兴初创企业，都积极开发应用于终端的人工智能芯片。

四、自动驾驶汽车实现完全自动化的时间表仍不明朗。目前，许多大型科技公司和初创企业正在大规模部署自动驾驶汽车领域，市场潜力巨大。2018 年本田在自动驾驶汽车领域投资 7.5 亿美元，软银投资 9 亿美元，Zoox 投资 5 亿美元；报告预计到 2025 年，该市场将达到 800 亿美元左右。尽管汽车制造商在全力以赴，但人们对这项技术的未来仍心存疑虑。最早采用全自动驾驶技术的物流配送目前仍处于试验和测试阶段，如必胜客一直在测试自动配送披萨等商品。

目前，人工智能技术已在面部识别、医学影像与诊断、电子商务搜索、自动驾驶汽车、网络安全等领域广泛应用，给社会生产生活带来了极大的舒适和便利。但在胶囊网络、下一代假肢等领域的应用尚处在试验阶段，随着技术的不断成熟，产品性能将得到极大提升，市场前景十分广阔。

附：《2019 年人工智能发展趋势》

（联系人及电话：李 鑫 15712824618）

撰稿：郝文姣　赛迪智库世界工业研究

人工智能的前景是什么？

2019 年人工智能发展趋势

AI Trends in 2019

作者：CBInsight 研究中心

来源：CBInsight 研究中心

编译：中国电子信息产业发展研究院

（赛迪研究院）

2019 年 3 月

目录

内容

NexTT 框架

2019 人 工 智 能 趋 势

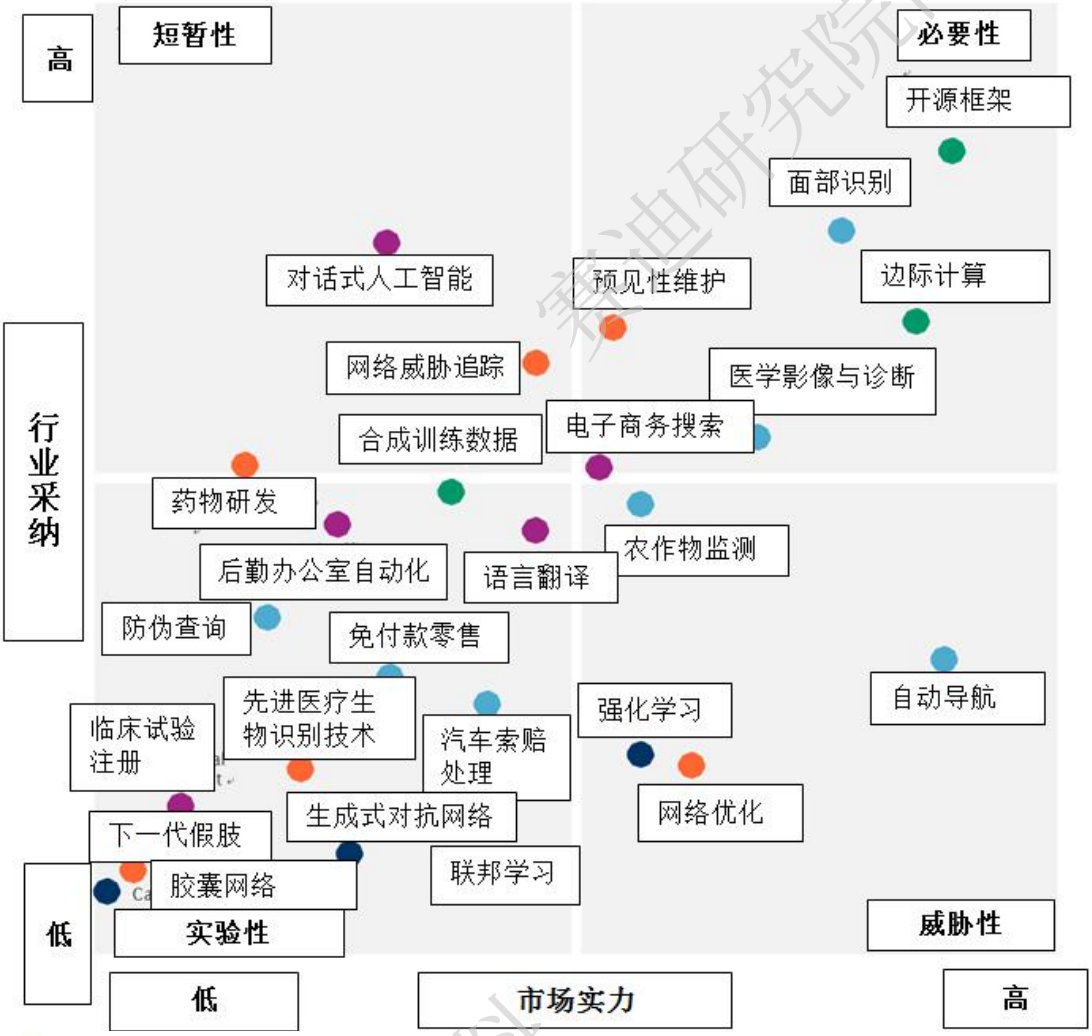

应用：计算机视觉

应用：自然语言处理/合成

应用：情报预测体系

结构

基础设施

NexTT 趋势

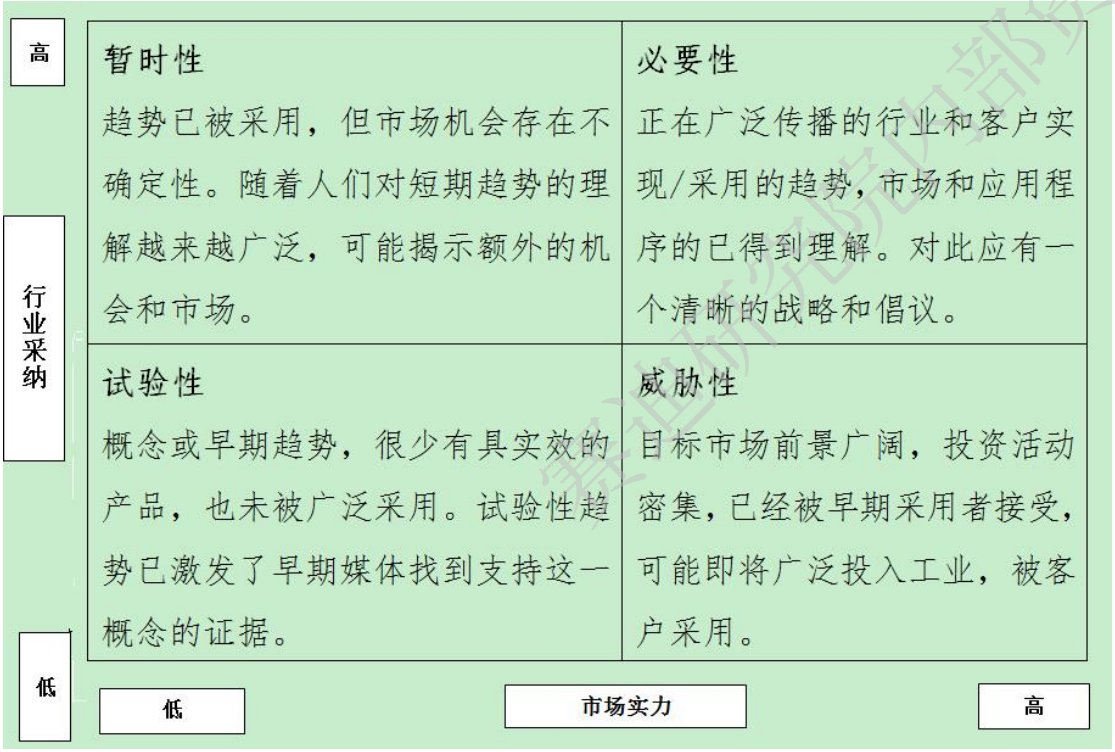

注：

我们采用 CB Insights 的 NExTT 框架评估这些趋势。

NExTT 框架帮助企业了解新兴趋势，并根据其对风险的适应程度为决策提供指导。NExTT 使用数据驱动信号来评估技术、产品和业务模型从概念到成熟再到广泛采纳的趋势。

NExTT 框架的 2 个维度：

行业采纳（y 轴）：包括该领域初创企业的势头、媒体关注、客户采纳（合作伙伴关系、客户、授权协议）。

市场实力（x 轴）：包括市场规模预测、投资者和资本的质量和数量、研发投资、收益记录评论、竞争强度、现有交易（并购和战略投资）。

NExTT 框架的 2 个维度

行业采纳（y 轴）
标志包括

该领域初创企业的势头

媒体关注

市场实力（X 轴）
标志包括

市场规模预测

收益记录评论

投资者和资本的质量和数量

竞争强度

研发投资

现有交易
（并购和战略投资）

必要性

1．开源框架

得益于开源软件，人工智能的门槛比以往任何时候都要低。

谷歌在 2015 年对外开放了 Tensor Flow 机器学习库。

人工智能的开源框架是双向的：每个人都可以访问人工智能，而谷歌这样的公司又受益于贡献者组成的社区，帮助其加快人工智能研究。

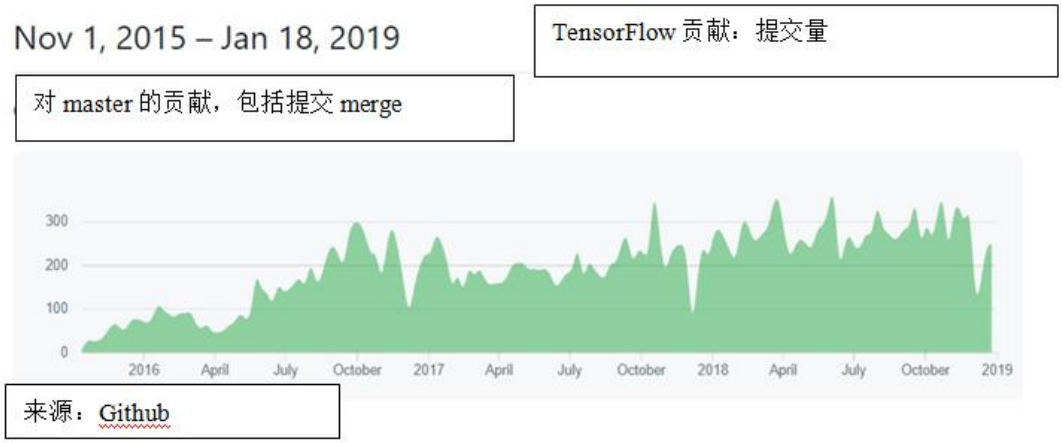

每个月都有数百名用户在 Git Hub（一个用户可以协作的软件开发平台）上为 Tensor Flow 做出贡献。

以下是一些使用 Tensor Flow 的公司，从可口可乐到 eBay，再到 Airbnb。

脸书（Facebook）于 2017 年发布了 Caffe2。在此之前，脸书与英伟达（Nvidia）、高通（Qualcomm）、英特尔（Intel）、微软（Microsoft）等公司的研究人员合作，旨在创建一个“轻量级、模块化的深度学习框架”，可以扩展到云以外的移动应用程序。

脸书当时也在运营 PyTorch，这是一个面向 Python 的开源机器学习平台。5 月 18 日，脸书将两种方法合二为一，“将 Caffe2 和 PyTorch 的优点结合起来，实现了快速原型设计和快速执行之间的平稳过渡。”

最近几个月，PyTorch 的贡献者数量增加了。

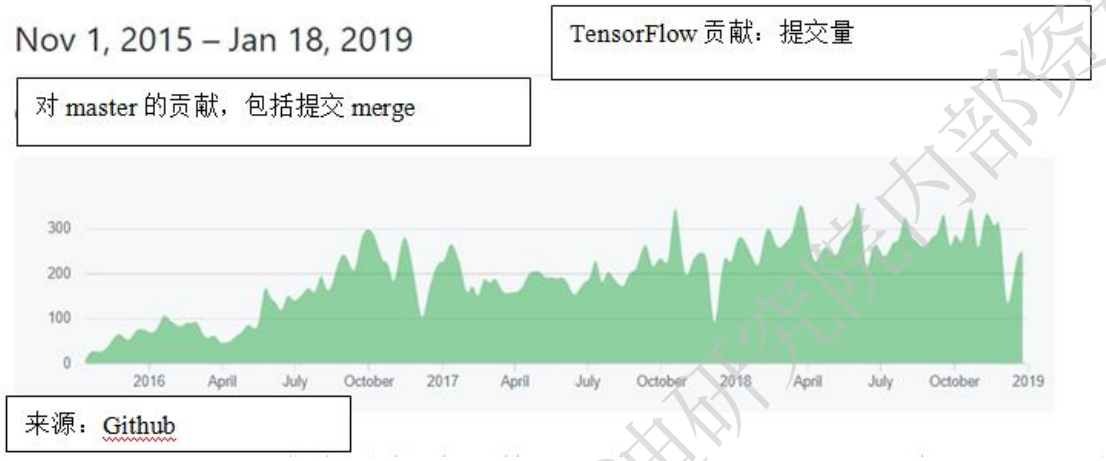

Theano 是蒙特利尔学习算法研究所（MILA）的另一个开源图书馆。9 月 17 日，业内领先的人工智能研究人员 Yoshua Bengio 宣布停止开发来自 MILA 的 Theano 工具，因为这些工具已经变得如此广泛。

“支持深度学习研究的软件生态系统一直在快速发展，现在已经达到一个健康的状态：开源软件是常态；框架种类各异，满足了从探索新思想到将其部署到生产中的各种需求；在一场激烈的竞争中，强大的行业参与者正在支持不同的软件堆栈。”

——Yoshua Bengio 对 MILA 的声明

现在有许多开放源码工具可供开发人员选择，包括 Keras、Microsoft Cognitive Toolkit 和 Apache MXNet。

2．终端人工智能

对实时决策的需求正将人工智能推向终端。

在智能手机、汽车甚至可穿戴设备等终端设备上运行人工智能算法，而不是与中央云或服务器通信，使设备能够在本地处理信息，并对情况做出更快的反应。

英伟达、高通和苹果以及许多新兴初创企业，都专注于

为处于“终端”的人工智能工作负载制造专门的芯片。

从消费类电子产品、电信到医疗成像，终端人工智能对每个主要行业都有影响。

例如，自动驾驶汽车必须对路上发生的事情做出实时响应，并在没有互联网连接的地区发挥作用。决策对时间十分敏感，延迟可能危及到生命。

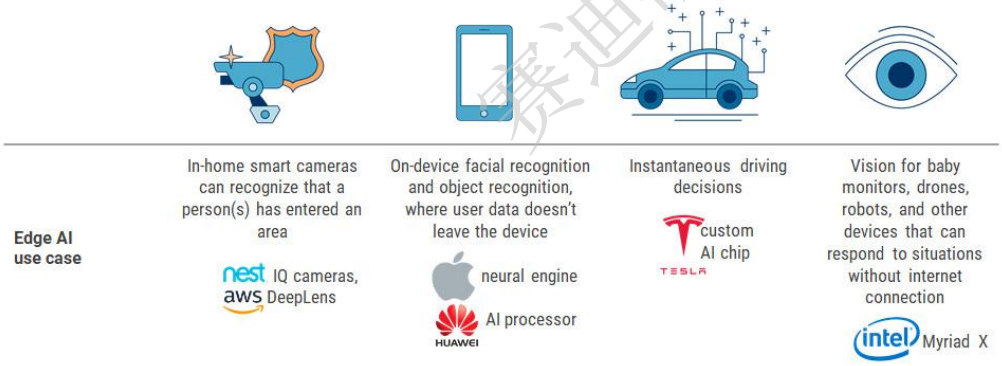

2017～2018 年，大型科技公司在终端人工智能领域实现了巨大飞跃。

2017 年，苹果发布了为 iPhone 8、iPhone 8 Plus 和 X 配备“神经引擎”的 A11 芯片，声称可以执行机器学习任务，速度高达每秒 6000 亿次运算，为 iPhone 提供了人脸识别等新功能，可以在设备上运行人脸识别来解锁手机。

高通在 2008 年第四季度推出了一项 1 亿美元的人工智能基金，投资于“那些与高通有共同愿景的初创企业，即让设备上的人工智能变得更强大、更广泛”。高通表示，此举与其 5G 愿景相辅相成。

作为许多数据中心的主导处理器，英特尔不得不在大规模收购中迎头赶上。英特尔发布了一款名为 Myriad X 的设备

内置视觉处理芯片（最初由 Movidius 开发，英特尔于 2016 年收购该公司）。在第 2018 年第四季度，英特尔推出了英特尔 NCS2（神经计算棒 2 代），这是由 MyriadX 视觉处理芯片驱动，运行在终端智能设备上的计算机视觉应用程序。这些终端包括智能家居设备和工业机器人。

CB Insights 的收益记录分析工具显示，提到终端人工智能的次数在 2018 年的部分时间里呈上升趋势。

来源：CB Insights 收入成绩单搜索引擎

微软表示，2018 年第三季度就引入了 100 项新的 Azure 功能，“既专注于现有的工作负载（如安全性），也专注于新的工作负载（如物联网和终端人工智能）。”

英伟达最近发布了 Jetson AGX Xavier 计算芯片，用于机器人和工业物联网领域的边缘计算应用。

虽然边缘计算的人工智能减少了延迟，但也有局限性。

与云不同，边缘有存储和处理限制，将出现更多的混合模型，允许智能边缘设备相互通信和与中央服务器通信。

3．面部识别

从解锁手机到登机，人脸识别正在成为主流。

在人脸识别方面，中国大力推进智能监控，再加上人工智能发展规划，已成为媒体关注的焦点。

在政府将人工智能技术用于监控的过程中，初创企业在为政府提供基础技术方面发挥了关键作用。在 CB Insights 平台上快速搜索中国人脸识别初创企业交易，就能看出对这项技术的需求。

像 SenseTime、Face++以及最近的 CloudWalk 这样的独角兽公司已在中国出现。（以下是我们关于中国监控活动的详细报告。）

根据 CB Insights 的专利分析工具，即使在美国，对这项技术的兴趣也在激增。

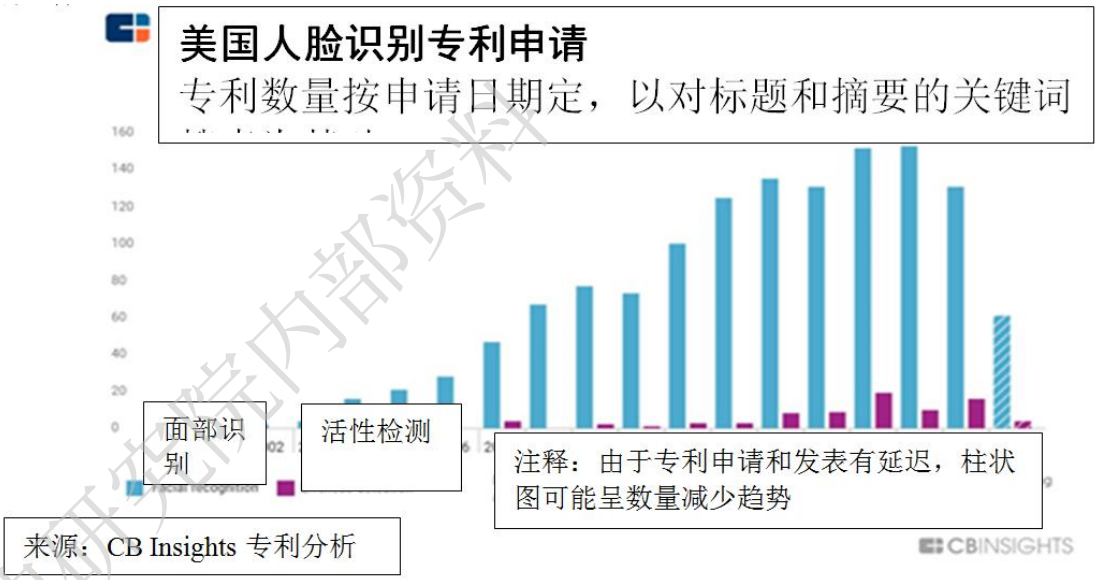

苹果公司在 iOS 10 中引入了基于面部识别的登录功能，

使这项技术普及到日常消费者。

亚马逊正在向执法机构出售其技术。

卡内基梅隆大学（Carnegie Mellon University）等学术机构也在研究帮助加强视频监控的技术，并获得了一项关于“幻觉面部特征”的专利。“幻化面部特征”是一种帮助执法机构识别蒙面嫌犯的方法，只需要捕捉到人脸的眼周区域，就可以重建完整的人脸。然后，面部识别可以用来将“幻化脸”与真实脸的图像进行比较，以发现相关性很强的人脸。

Source: USPTO

但这项技术并非没有瑕疵。据报道，亚马逊因为将一些国会议员误认为罪犯而上了新闻。

西雅图一所学校校外的智能摄像头很容易被《华尔街日报》（WSJ）的一名记者蒙骗。当“微笑解锁”功能暂时被禁用时，这位记者用一张校长的照片进入了学校。

“微笑解锁”和其他类似的“活性检测”方法提供了额外认证。

亚马逊获得的一项探索额外安全防护的专利要求用户执行某些动作，如“微笑、眨眼或倾斜头部”。然后，将这些操作与“红外图像信息、热成像数据或其他此类信息”相结合，以实现更具说服力的身份验证。

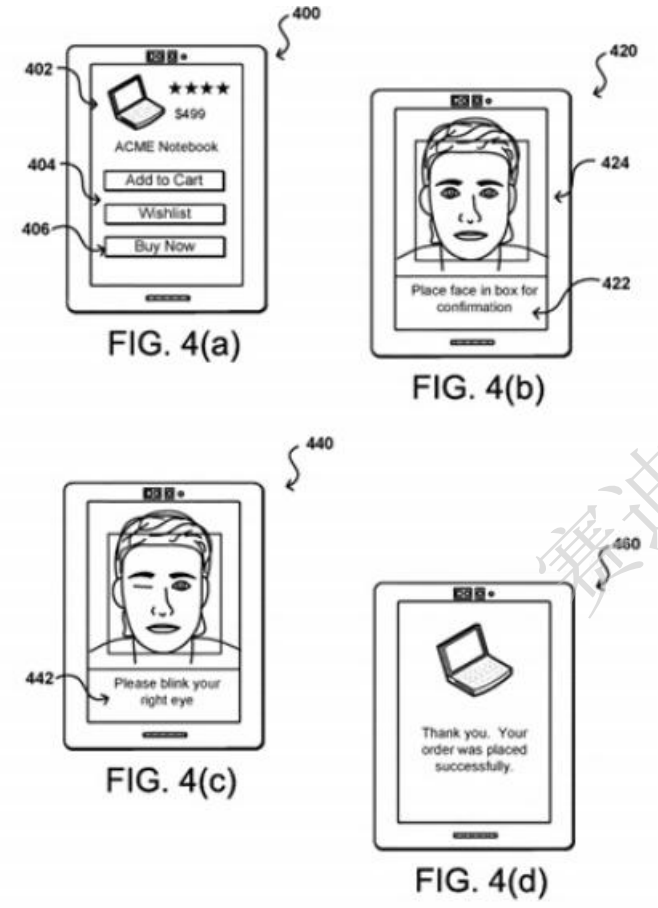

Source: USPTO

在安全、零售和消费电子领域，早期的商业应用正在腾飞，面部识别正迅速成为生物特征认证的主导形式。

4．医学影像与诊断

美国食品药品监督管理局（FDA）正在为人工智能作为医疗设备开绿灯。

2018 年 4 月，FDA 批准了一款人工智能软件，该软件无需来自专家的二次意见，即可筛查糖尿病视网膜病变患者。

该产品被授予“突破性设备设计”奖励，以加快其市场化进程。

IDx-DR 软件识别出“轻度以上糖尿病性视网膜病变”患者的正确率为 87.4%，识别是否患糖尿病性视网膜病变患

者的正确率为 89.%。近几个月来，IDx 是 FDA 批准用于临床商业应用的众多人工智能软件产品之一。

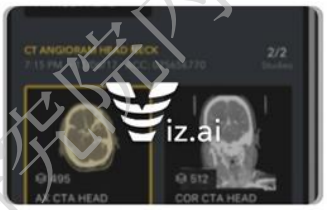

糖尿病视网膜病变肝和肺 AI 期病变　　CT 扫描分析中风

FDA 批准 Viz LVO （来自创投企业 Viz.ai 的产品）对 CT 扫描进行分析，并向医疗服务提供者通知患者可能出现中风的情况。FDA 批准后，Viz.ai 完成了谷歌风投（GoogleVentures）和凯鹏华盈（Kleiner Perkins Caufield & Byers）价值 2100 万美元的首轮融资。

FDA 还批准了由通用电气（GE）风投支持的初创企业 Arterys 研究肿瘤人工智能管理软件。该软件最初专注于发现肺部和肝脏病变。

快速审批为人工智能成像与诊断公司开辟了新的商业途径。自 2014 年以来，在 149 交易中，已有 80 多家公司通过了股权融资。

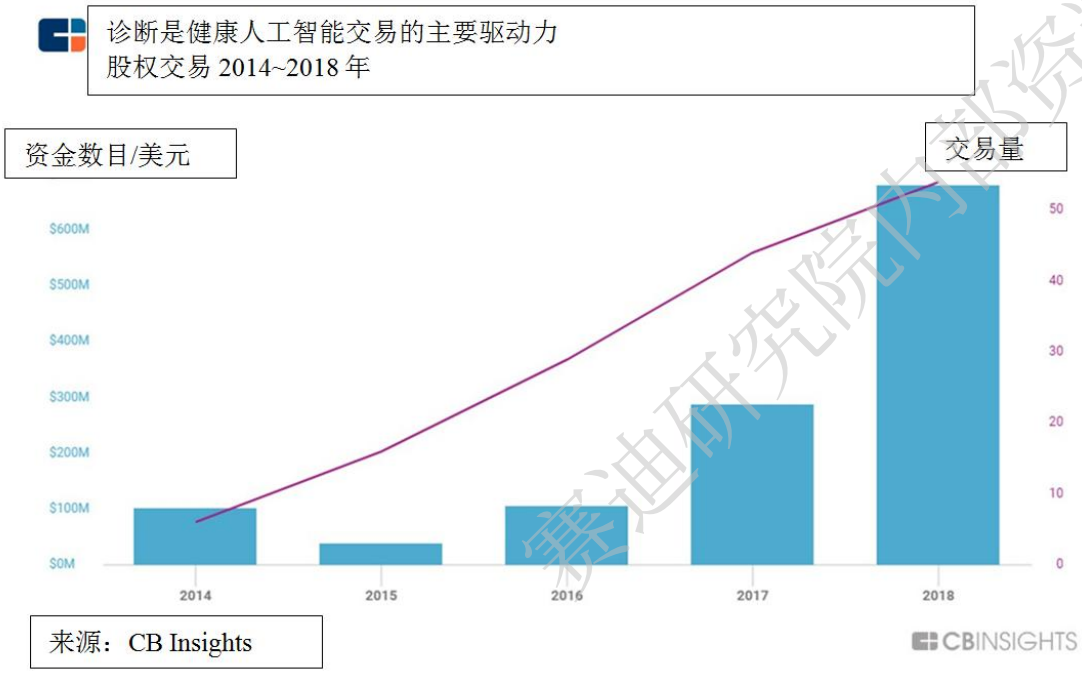

在消费者方面，智能手机的普及和图像识别技术的进步正在把手机变成强大的家庭诊断工具。初创企业 Health.io 的第一款产品是 Dip.io。Dip.io 使用传统的尿液分析试纸来监测一系列的尿路感染。使用者用智能手机拍下试纸的照片，计算机视觉算法根据不同的光照条件和相机质量对结果进行校正，从而检测感染和妊娠相关并发症。

Dip.io 已经在欧洲和以色列上市，并获得了 FDA 的批准。

除此之外，许多医疗成像与诊断服务平台正在与获 FDA 批准的家庭监控设备集成，在出现异常时向医生发出警报。

5．预见性维护

从制造商到设备保险公司，AI-IIoT 可以为现有企业在意外故障方面节省数百万美元。

现场和工厂设备产生大量的数据，但意外的设备故障是导致生产停工的主要原因之一。

通用电气最近对 450 家现场服务和 IT 决策者进行的一项

调查发现，70%的公司不知道设备何时需要升级或维护，而意外的停工可能会让公司每小时损失 25 万美元。

预测设备或单个部件何时会出现故障，不仅有利于制造商，也有利于资产保险公司。

在预见性维护中，传感器和智能相机从机器上收集连续的数据流，比如温度和压力。实时数据生成的数量和格式的多样性使得机器学习成为 IIoT 不可分割的组成部分。随着时间的推移，这些算法可以预测即将出现的故障。

随着工业传感器成本的下降，机器学习算法的进步，以及边缘计算的推动，使得预见性维护更加广泛可用。

与这一领域相关的一个领先指标是这里的大型科技公司和初创企业的绝对数量。

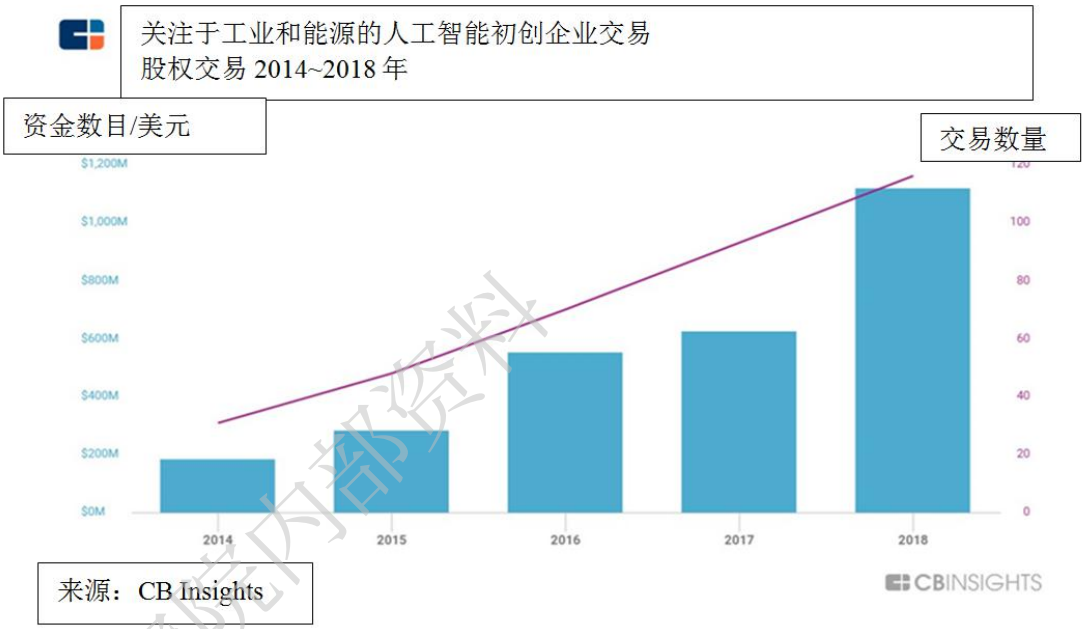

专注于工业和能源领域的人工智能公司的交易正在增加，其中包括面向 IIoT 的 ML-as-a-service 平台。更新兴的初创企业正在与 C3 物联网和 Uptake Technologies 等独角兽公

司竞争。

2016 年，GE 风投是该领域的活跃投资者，支持的公司包括 Foghorn Systems、Sight Machine、Maana 和 Bit Stew Systems（后来被 GE 风投收购）。通用电气是 IIoT 的主要参与者，拥有 Predix 分析平台。

竞争对手包括西门子和 SAP，已经推出了各自的 IIoT 产品（Mindsphere 和 Hana）。

印度塔塔咨询公司宣布，将为能源公用事业公司推出预见性维护和基于人工智能的解决方案。塔塔声称，其“数字孪生”技术的初期版本将以数字格式复制地面业务或实物资产，以实现监测，可以帮一个发电厂每年节省约 150 万美元/千兆瓦。

甚至像微软这样的大型科技公司也在扩展云计算和边缘分析解决方案，包括预见性维护。

6．电子商务搜索

对搜索词上下文语境的理解正在走出“实验阶段”，但要被广泛采用还有很长的路要走。

自 2002 年以来，亚马逊已经申请了 35 项与“搜索结果”相关的美国专利。亚马逊有一个独家子公司 A9，专注于亚马逊的产品和视觉搜索。A9 在美国有近 400 项专利申请（并非都与搜索优化有关）。

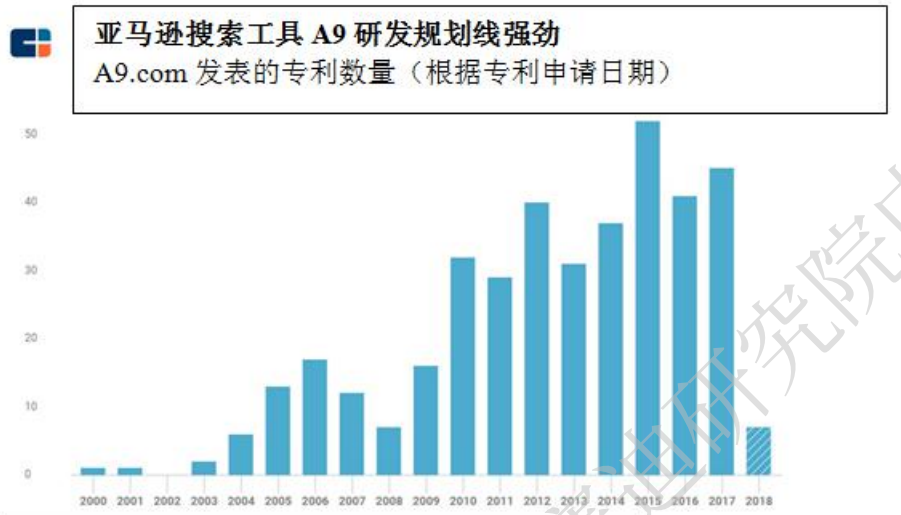

注释：由于专利申请和发表有延迟，柱状图可能呈数量减少趋势
来源：CB Insights 专利分析

一些与搜索相关的专利包括使用卷积神经网络“确定一组项目，使其图像与查询图像显示视觉相似性……”，以及使用机器学习分析图像的视觉特征，并基于这些特征构建搜索查询。

亚马逊为搜索部门专门招聘了 150 多个职位，包括自然语言理解、混沌工程和机器学习等。

但在零售商中，亚马逊在电子商务搜索领域的运营和研发规模是个例外。

很少有零售商在财报电话会议上讨论与人工智能相关的策略，许多零售商也没有扩大或优化其电子商务业务。

但最早这样做的品牌之一是 eBay。

eBay 首次提到“机器学习”是在 2015 年第三季度的财报电话会议。当时，eBay 刚刚开始强制要求卖家写产品描述，并使用机器学习来处理这些数据，以便在目录中找到类似的产品。

使用适当的元数据描述网页站点上的产品是使用电子商务搜索显示相关搜索结果的起点。

但是仅仅描述和索引是不够的。许多用户用自然语言搜索产品（比如“一件没有纽扣的洋红色衬衫”），或者不知如何描述他们正在寻找的东西。

这使得电子商务搜索的自然语言成为一个挑战。

处于早期的 SaaS 初创公司正在兴起，并开始向第三方零售商销售搜索技术。

图片搜索初创公司 ViSenze 与优衣库（Uniqlo）、Myntra 和日本电子商务巨头乐天（Rakuten）等客户合作。ViSenze 允许店内顾客为他们喜欢的商品拍照，然后上传照片，在网上找到确切的商品。

该公司在加州和新加坡设有办事处，并在 2016 年从包括乐天风投部门在内的投资者筹集 1050 万美元的 B 轮融资。

该公司于 2017 年成为联合利华创想+，允许东南亚的初创企业用自己的品牌开展试点项目。

另一家为在线搜索推荐做开发人工智能的初创公司是以色列的 Twiggle。阿里巴巴支持这家公司开发一种基于现有电子商务搜索引擎的语义 API（应用程序编程接口），以对来自买家非常具体的搜索做出相应。Twiggle 在 2017 年的 B 轮融资中筹集到 1500 万美元，并于 2018 年进入即插即用加速器领域。

试验性

1．胶囊网络

如今，大多数人工智能应用程序都采用深度学习。胶囊网络领域为深度学习开辟了新应用。

谷歌的杰弗里•辛顿（Geoffrey Hinton）是深度学习领域的先驱研究员，他早在 2011 年就在一篇论文中引入了“胶囊”的新概念，并认为“目前的图像识别方法效果不佳，使用的方法不聪明，无法令人满意。”

辛顿提到的这些“目前方法”包括当今深度学习中最流行的神经网络架构之一，即卷积神经网络（CNN）。CNN 在图像识别方面的应用尤其迅速。尽管 CNNs 取得了成功，它也有不足之处（详见下文）。

辛顿在 2017～2018 年发表了两篇关于“胶囊网络”（Caps Net）的论文，这是一种新的架构，有望在多个方面超越 CNNs。

如果不深入研究其不足之处，CNNs 在精确的空间关系方面就会失败。看看下面这张脸，虽然嘴的相对位置与其他面部特征不相符，但 CNN 仍会将其识别为人脸。

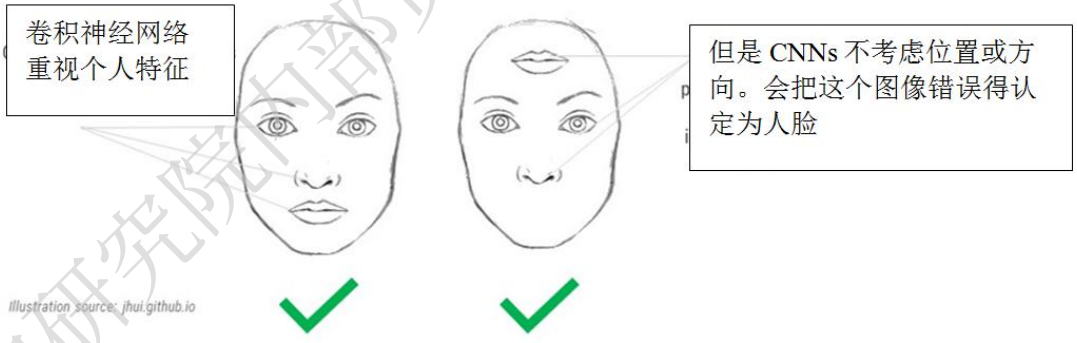

尽管有一些方法可以缓解上述问题，但 CNNs 的另一个

主要问题是无法理解新观点。

“现在卷积神经网络已经成为物体识别的主流方法，讨论是否存在导致其消亡的指数级低效能，是有意义的问题。卷积神经网络在概括并形成新颖观点方面存在的困难可能是个好答案。”

——关于胶囊之间动态路由的论文

例如，Caps Net 在识别第一和第二行玩具的图像时做得更好，它们属于同一个对象，只是从不同的角度或角度拍摄的。CNNs 需要更大的训练数据集来识别每个方向。

（上面的图片来自一个叫做 smallNORB 的数据库，其中包含了 50 个玩具的灰色图像，每个都分属于 5 个类别中的 1 个：四足动物、人像、飞机、卡车和汽车。辛顿的论文发现，与其他算法方法相比，在该数据集上进行测试时，Caps Nets 的错误率降低了 45%。）

辛顿声称，经过复杂的对抗攻击测试（通过篡改图像来混淆算法）后，胶囊网络性能优于卷积神经网络。

黑客可以引入一些小的变化来愚弄 CNN。谷歌和 OpenAI 的研究人员通过几个例子证明了这一点。

Caps Net 测试的一个比较典型的例子，是来自谷歌的 Ian

Goodfellow 等人 2015 年的一篇论文。如下图所示，一个人眼不易察觉的微小变化意味着，这幅图像产生了一个神经网络，将熊猫识别为长臂猿，且可信度很高。

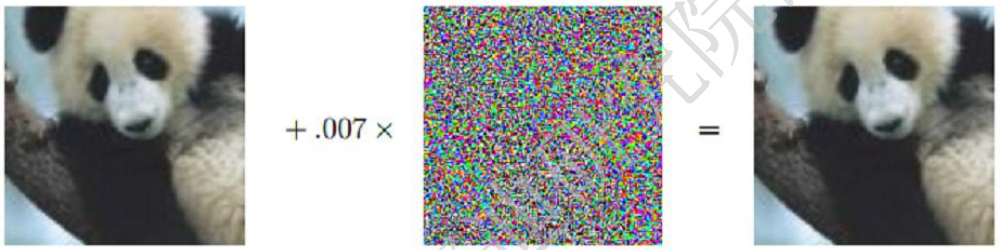

$\boldsymbol{x}$
"panda"
57.7% confidence

$\text{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, y))$
"nematode"
8.2% confidence

$\boldsymbol{x} + \epsilon\text{sign}(\nabla_{\boldsymbol{x}} J(\boldsymbol{\theta}, \boldsymbol{x}, y))$
"gibbon"
99.3 % confidence

胶囊网络的研究还处于起步阶段，但可能会挑战目前最先进的图像识别方法。

2．下一代假肢

处于早期的研究正在兴起，结合生物学、物理学和机器学习来解决假肢中最困难的问题之一：灵活性。

DARPA 已经在其高级假肢项目上投入数百万美元，该项目始于 2006 年，目的是帮助受伤的退伍军人。但这是一个复杂的问题。

例如，让截肢者能够在假肢手臂上移动自己的手指，解码任意动作背后的大脑和肌肉信号，并将其转化为机器人控制，这都需要跨学科的方法。

正如 2018 年 Megan Molteni 在《连线》杂志的一篇文章中解释的那样，以弹钢琴为例，经过反复练习，和弦变成了“肌肉记忆”，但假肢的原理并非如此。

最近，研究人员开始使用机器学习来解码来自身体传感

器的信号，并将它们转换成移动假肢设备的命令。

约翰霍普金斯大学的应用物理实验室正在进行一个关于假肢神经接口的项目，使用“神经解码算法”来实现这一点。2018 年 6 月，德国和伦敦帝国理工学院（Imperial College London）的研究人员利用机器学习来解码截肢者残肢发出的信号，并驱动一台电脑来控制机械手。关于“脑机接口”的研究发表在《科学机器人》上。

其他一些论文探索了中间的解决方案，比如使用肌电信号（残肢附近肌肉的电活动）来激活相机，以及运行计算机视觉算法估算在他们面前的物体的抓取类型及大小。

“人工智能假肢挑战赛”是 2018 年 NeurIPS（一个领先的年度机器学习大会）的竞赛项目之一，进一步凸显了人工智能社区对该领域的兴趣。

2018 年的挑战是使用强化学习来预测假肢的性能（关于强化学习的更多信息见本报告的以下部分）。研究人员使用一种名为 OpenSim 的开源软件来模拟人体运动。

2018 年的重点是“学习跑步”，有 442 名参与者试图教人工智能如何跑步，赞助商包括 AWS，英伟达和丰田。

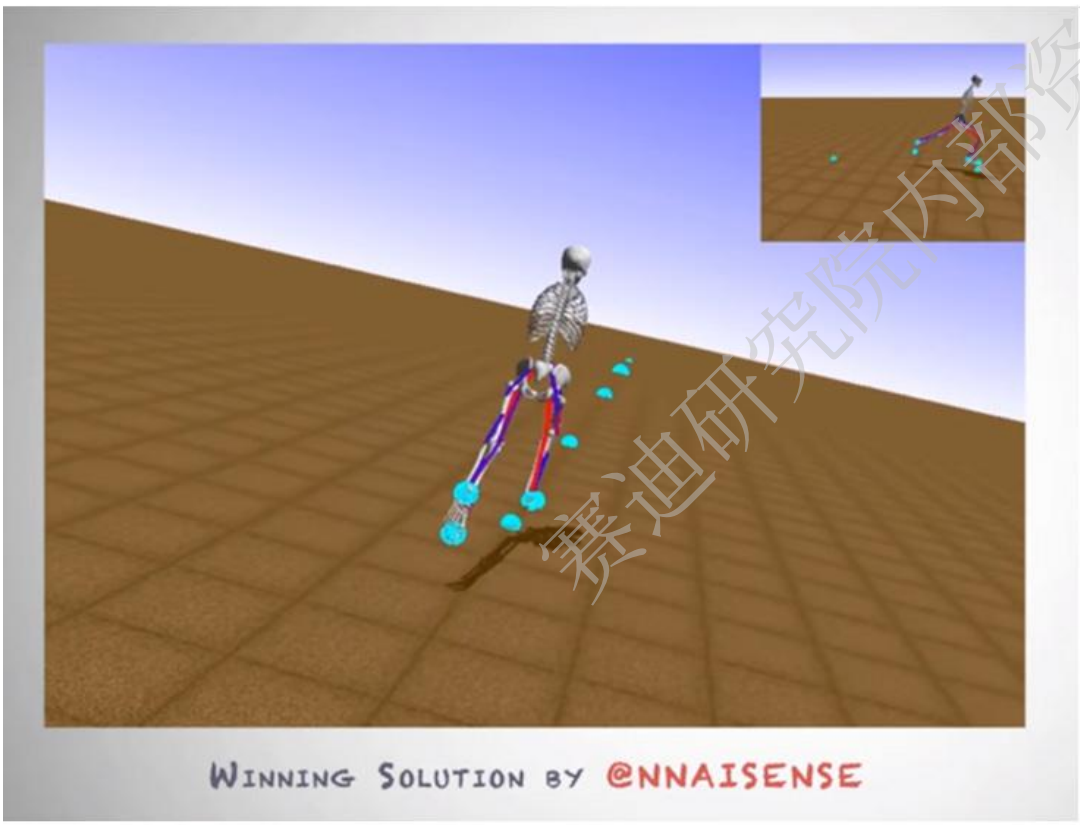

Source: http://osim-rl.stanford.edu/

3. 临床试验注册

临床试验的最大瓶颈之一是招募合适的患者。苹果或许能够解决这个问题。

互操作性，也就是跨机构和软件系统轻松共享信息的能力，是医疗保健领域最大的问题之一，尽管业界一直在努力实现健康记录的数字化。

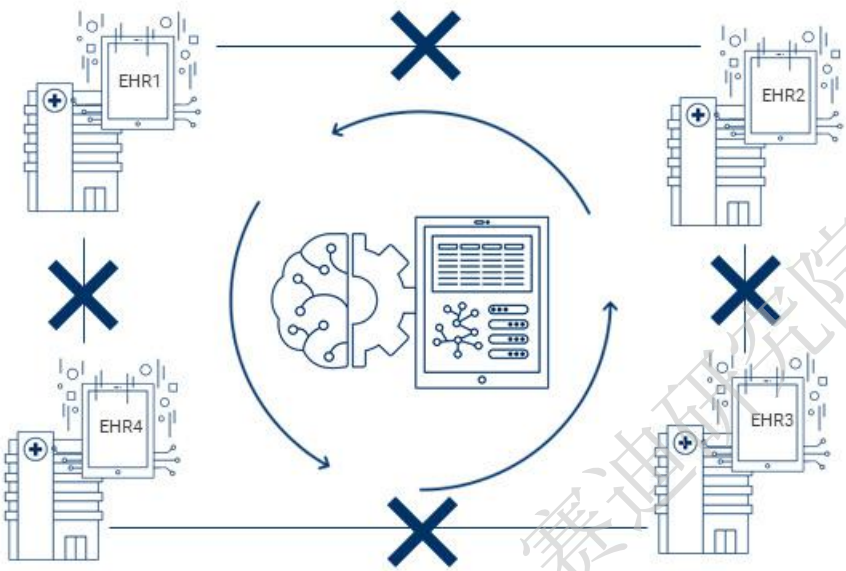

对临床研究团队和患者来说，在临床试验中，将正确的试验与正确的患者进行匹配是一个耗时且具有挑战性的过程。

在此背景下，仅在美国就有超过 18000 项临床研究正在招募患者。

如果医生知道正在进行试验，病人就可以偶尔从医生那里得到试验建议。

否则，通过临床试验综合联邦数据库 ClinicalTrials.Gov 获得数据的事项就落到病人身上。

理想的人工智能解决方案应该是利用人工智能软件从患者的病历中提取相关信息，与正在进行的试验进行比较，并建议进行匹配的研究。

很少有初创公司直接在临床试验领域与客户合作。小型初创公司简化临床试验的最大障碍是，技术相对较新，行业适应速度较慢。

然而，苹果等科技巨头已经成功地为其医疗保健计划找到了合作伙伴。

苹果正在改变医疗保健领域的数据流动方式，并为人工智能开辟了新的可能性，尤其是围绕临床研究人员招募和监控患者的方式。

2015 年以来，苹果推出了两个开源框架——ResearchKit 和 CareKit，帮助临床试验招募患者，并远程监控其健康状况。

这些框架允许研究人员和开发人员创建医疗应用程序，监控人们的日常生活，消除了参与试验的地理障碍。

例如，近 1 万人使用 mPower 应用程序，该程序提供手指敲击和步态分析等练习，用于研究帕金森病患者，这些患者们同意与更广泛的研究社区共享数据。

杜克大学（Duke University）的研究人员开发了一款名为“自闭症与超越”（Autism & Beyond）的应用程序，使用 iPhone 的前置摄像头和面部识别算法来筛查患有自闭症的儿童。

苹果还与 Cerner 和 Epic 等流行的 EHR 供应商合作解决互操作性问题。

2018 年 1 月，苹果公司宣布，iPhone 用户可以通过他们的 iPhone 健康应用程序访问参与机构的所有电子健康记录。

该功能名为“健康记录”（Health Records），是人工智能医疗初创企业 Gliimpse 所做努力的延伸，Gliimpse 于 2016 年被苹果（Apple）收购。

在一个易于操作的界面，用户可以找到所需的所有信息，

包括过敏、身体状况、免疫、实验室结果、药物、程序和生命体征。

2018 年 6 月，苹果为开发者推出了健康记录 API。用户现在可以选择与第三方应用程序和医学研究人员共享数据，为疾病管理和生活方式监测开辟了新的机会。

在使用人工智能和机器学习进行早期诊断、招募合适的患者、甚至在药物设计中推动决策方面，仍拥有无限可能。

4. 生成式对抗网络

两个神经网络试图在生成逼真图像方面一较高下。

你能认出这些图片中哪些是假的吗?

答案是以上所有图像。这些高度逼真的图像都是由生成式对抗网络（GANs）创建的。

（注意：右下角的图像表示“类泄漏”，算法可能混淆了狗和球的属性，并创建了一个“两不像”）

2014 年，谷歌研究员伊恩•古德费洛提出了“生成式对抗网络”的概念，引入了“AI 和 AI 相对抗的想法”。“有两种神经网络：一种是生成器，生成一张假图（比如一只狗）；另一种是识别器，将结果与真实图像进行比较并反馈给生成器，告诉它离复制真实图像的距离有多近。”

这就形成了两个神经网络之间不断的反馈回路，它们都想要以智取胜，压倒对方。

上面的图片来自于谷歌 Deep Mind 的实习生 Andrew Brock 于 2008 年 9 月 18 日发表的一篇论文，该论文与 Deep Mind 的其他研究人员一起发表。他们在一个非常大的数据集上训练 GANs 来创建“BigGANs”。

Brock 及其团队在处理 BigGANs 时遇到的一个挑战是：例如，一只蜘蛛有“很多条腿”。但是“很多”是多少呢?

然而，扩展像 GANs 这样大型项目的主要挑战是计算能力。下面是来自 Fast Company 的一段摘录，粗略估计了这项研究的计算能力：

这些试验对环境也有影响。Brock 使用 512 张谷歌张量处理单元（TPU）生成 512 像素的图像，实验通常运行 24 到 48 小时。如果每个 TPU 在一个小时的计算中使用大约 200 瓦，那么在 Brock 的 512 像素实验中，一个单独的实验可以使用 2450 到 4915 千瓦小时。**这相当于美国普通家庭在 6 个月内使用的电量。**

为了使 GANs 能够扩展，人工智能的硬件必须并行扩展。

Brock 的论文并不是最近几个月发表的唯一与 GAN 相关的论文。

英国兰卡斯特大学（Lancaster University）、中国西北大学（Northwest University）和中国北京大学（Peking University）的研究人员利用 GANs 开发了一个验证码求解器。

本文证明了 GANs 使用桌面 GPU 可以在 0.05 秒内破解基于文本的验证码，与之前的方法相比成功率相对较高。

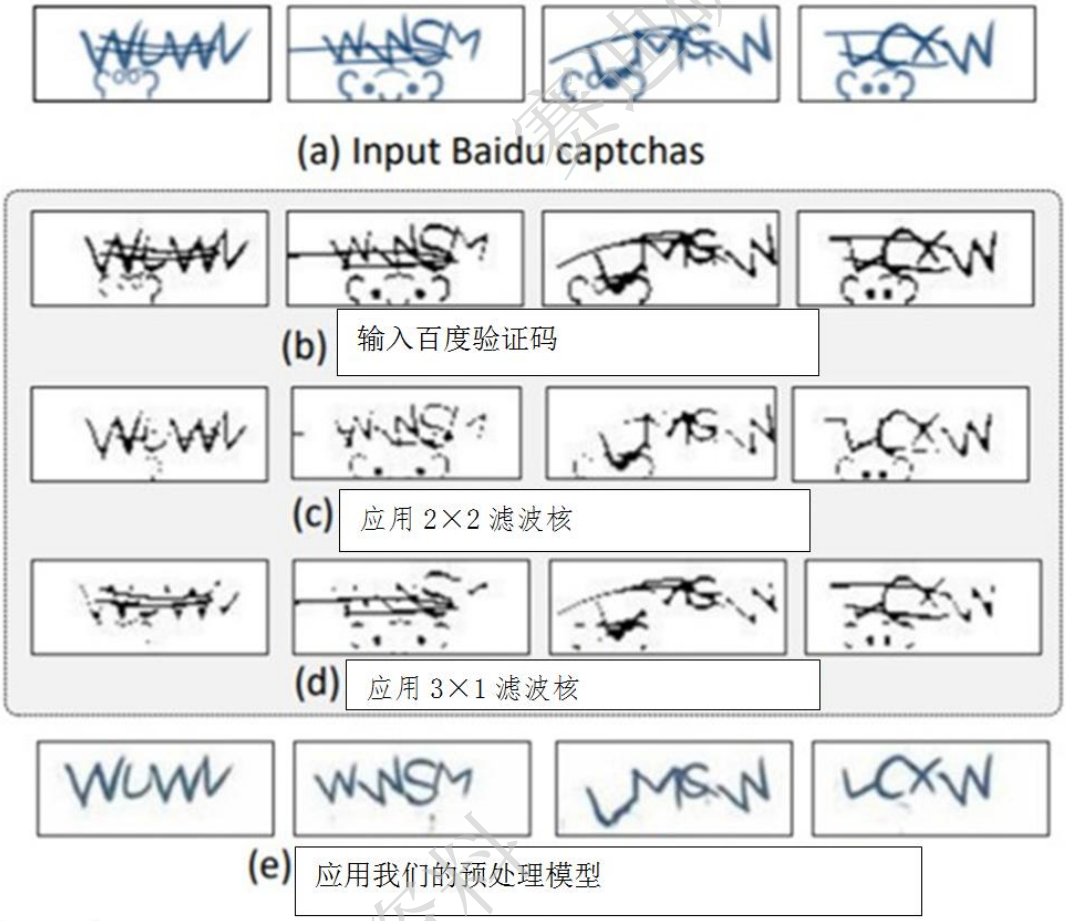

来源：lancaster.ac.uk/staff/wangz3/publications/ccs18.pdf

CMU 的研究人员在这个“深度伪造”视频的迭代中使用 GANs 进行“面对面”转换。在下面的例子中，John Oliver 变成了 Stephen Colbert：

Source: http://www.cs.cmu.edu/~aayushb/Recycle-GAN/

华沙科技大学的研究人员开发了一个 ComixGAN 框架，可以使用 GAN 将视频转换成漫画。

图像 8：不同 comixification 方法的结果。列：a 内容图片, b 原始 CartoonGAN 模型(Hayao)，c 原始 CartoonGAN 模型(Hosoda)，dComixGAN

艺术品拍卖行佳士得（Christie's）以 43.25 万美元的天价售出了史上由 GANs 首次创作的油画。

Portrait of Edmond Belamy, 2018, created by GAN (Generative Adversarial Network). Sold for $432,500 on 25 October at Christie's in New York. Image © Obvious

在最近一篇关于 GANs 的论文中，Nvidia 的研究人员使用了一种“基于样式的生成器”来创建超高清图像。

Source: https://arxiv.org/pdf/1812.04948.pdf

GANs 不只是用于好玩的实验。这种做法也有严重的影响，包括伪造政治视频和变相色情。《华尔街日报》（The Wall Street Journal）已经在训练研究人员识别深度假视频。

随着研究规模的扩大，GANs 将改变新闻、媒体、艺术甚至网络安全的未来。GANs 已经在改变我们训练人工智能算法的方式（更多内容将在 “联合学习”中介绍）。

5．联合学习

新方法旨在保护隐私，同时用敏感用户数据训练人工智能。

我们与智能手机和平板电脑的日常互动——从我们在短信中使用的词语选择到我们对照片的反应——产生了大

量的数据。

使用独特的本地数据集训练人工智能算法，可以极大地提高其性能，比如更准确地预测要输入键盘的下一个单词。

正如谷歌的研究人员在 2017 年的一篇论文中所解释的那样，“聊天和短信中使用的语言通常与标准语料库大不相同，如维基百科和其他网络文档，人们在手机上拍的照片很可能与 Flickr 上的普通照片大不相同。”

但这些用户数据属于个人，是敏感隐私。

谷歌的联合学习方法旨在使用这个丰富的数据集，但同时保护敏感数据。

简而言之，保存在手机上的数据不会被发送到中央云服务器或存储在中央云服务器中。云服务器将算法的最新版本——所谓算法的“全局状态”——发送给随机选择的用户设备。

手机会根据本地化数据对模型进行改进和更新。只有这个更新（以及来自其他用户的更新）被发送回云，改进“全局状态”，这一过程会重复进行。

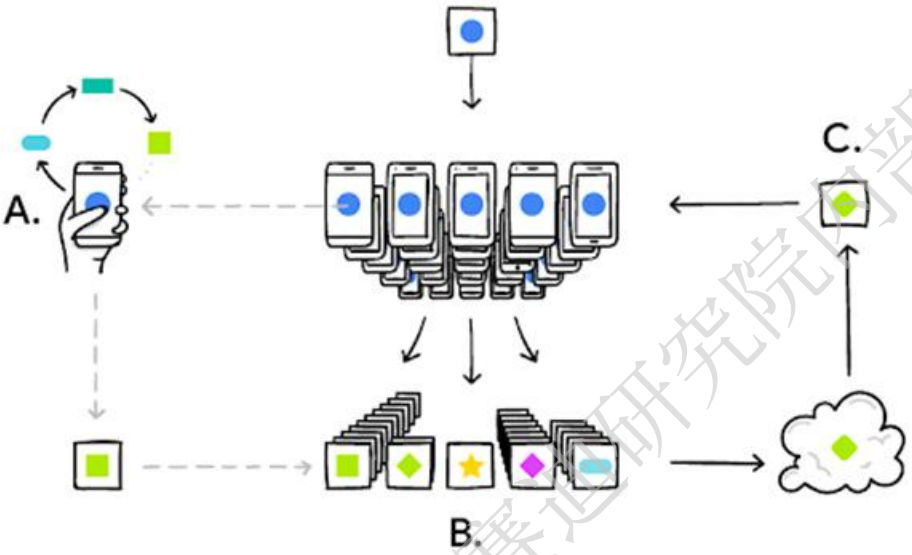

Source: Google AI blog

谷歌正在其名为 Gboard 的安卓键盘上测试联合学习。

注意，聚合来自每个节点的个体更新聚合机制并不新鲜。有些算法已经做到了这一点。

但与其他分布式算法不同，联合学习方法考虑了数据集的两个重要特征：

•非独立同分布（Non-IID）：根据每个人对设备的使用情况，每个手机（或其他设备）上生成的数据是唯一的。因此，这些数据集并不是“独立且同分布的（IID）”——这是其他分布式算法统计推断后的常见假设，但不反映实际的现实情况。

•不平衡：一些用户比其他人更积极地使用应用程序，自然会产生更多的数据。因此，举例来说，每部手机都会有不同数量的培训数据。

Firefox 测试了联合学习，以对用户开始输入 URL 栏时出现的建议进行排序，并将其称为“大型软件项目中最早运行‘联合学习’的举措之一”。

联合学习的另一个应用是，谷歌风投支持的专注于药物研发的人工智能初创企业 OWKIN 正在使用这种方法来保护敏感的患者数据。投资者 Otium Venture 表示，该模型允许不同的癌症治疗中心在不需要患者数据的情况下进行合作。

6．先进医疗生物识别技术

利用神经网络，研究人员开始研究和测量以前难以量化的非典型风险因素。

利用神经网络分析视网膜图像和语音模式，可能有助于识别患心脏病的风险。

根据 2019 年发表在《自然》（Nature）杂志上的一篇论文，谷歌研究人员使用了一个在视网膜图像上受训的神经网络，来寻找心血管风险因素。

研究发现，不仅可以通过视网膜图像识别年龄、性别和吸烟模式等危险因素，还可以“在一定程度上精确量化，这是以前从未报道过的。”

同样，梅奥诊所与以色列初创企业 Beyond Verbal 合作，分析声音的声学特征，以发现冠状动脉疾病（CAD）患者不同的声音特征。研究发现，当受试者描述情感体验时，两种声音特征与 CAD 密切相关。

初创公司 Cardigram 最近的研究表明，利用深度学习，“由糖尿病驱动的心率变异性变化可以通过消费者的自备可穿戴心率传感器检测出来”。一种算法方法显示从心率检测糖尿病的准确率为 85%。

一个在未来更具前景的用例是医疗生物识别的被动监

视。

2018 年 1 月，谷歌发布了一项专利，该专利专注于从一个人的肤色或皮肤位移来分析心血管功能。

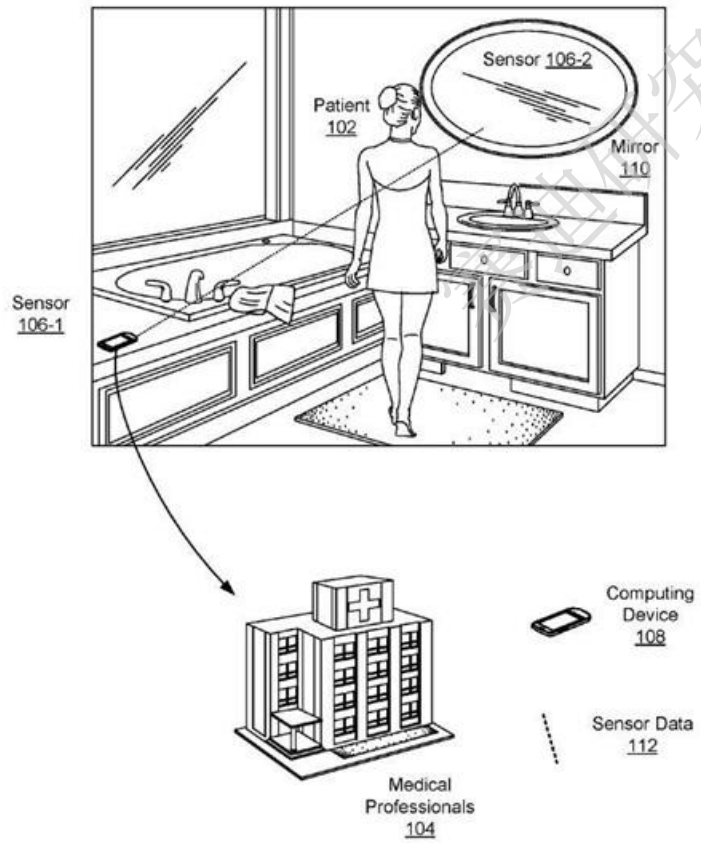

Source: USPTO

这些传感器甚至可能被放置在病人浴室的“感应环境”中（根据专利说明）。

例如，通过识别手腕和脸颊的皮肤颜色变化，并“比较（测量）时间和这些区域之间的距离”，系统可以计算出“脉冲波速（PWV）”。

速度信息可以用来确定心脏健康指标，如动脉僵硬度或血压。

该专利称:“机器学习可以应用于为患者创建一个特定的模型，用 PWV 估计血压。”

亚马逊在 2014 年申请了一项类似的被动监测专利， 并于 2017 年获得批准，可以将人脸特征识别（使用神经网络或其他算法）与心率分析结合起来。

例如，算法可以跟踪人脸两个区域的颜色变化，比如眼睛和脸颊附近的区域，用于心率检测计算。

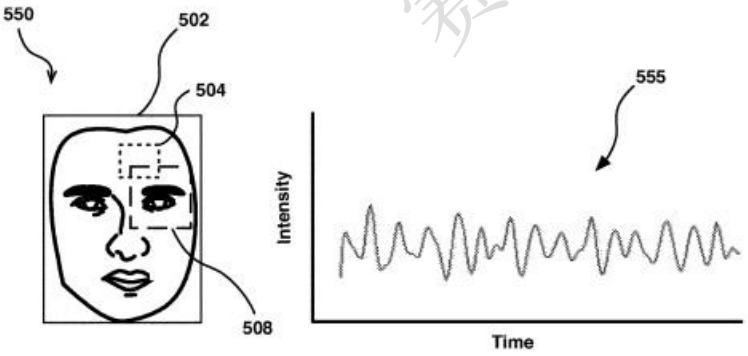

FIG. 5E

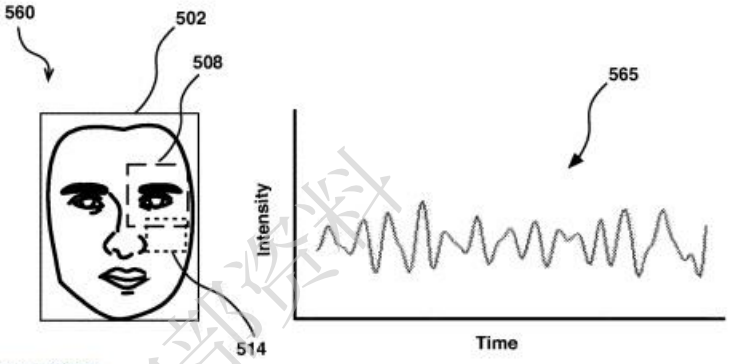

Source: USPTO

人工智能发现模式的能力将继续为新诊断方法识别以前未知的风险因素铺平道路。

7. 汽车索赔处理

保险公司和初创公司开始使用人工智能来计算车主的

“风险评分”，分析事故场景的图像，并监控司机的行为。

阿里巴巴旗下的蚂蚁金服在其“事故处理系统”中使用深度学习算法进行图像处理。

目前，车主或司机把车交给“理赔员”，理赔员负责检查车辆的损坏情况，并记录下详细信息，然后将这些信息发送给汽车保险公司。

随着图像处理技术的进步，人们现在可以给车拍照并上传到蚂蚁金服上。然后，神经网络分析图像并自动进行损伤评估。

蚂蚁金服采取的另一种方法是创建驾驶员影响汽车保险实际定价的风险概况。

“大数据和人工智能等技术的发展，使保险公司能够进一步发挥杠杆作用对消费者数据进行分析，并对车主可能暴露的风险进行分析。因此，汽车保险的风险因素可以从“以汽车为导向”转变为“人车结合”。

——阿里巴巴云博客

阿里巴巴推出了名为“汽车保险积分”（Auto Insurance Points）的系统，利用机器学习，根据信用记录、消费习惯和驾驶习惯等因素，计算车主的风险得分。

规模较小的初创企业也在涉足保险和理赔业务，但采取了不同的方式。例如，Nexar 鼓励司机将智能手机当作行车记录仪，并将视频上传到 Nexar 的应用程序中。作为回报，车主支付保险费可以获得折扣。

该应用程序使用计算机视觉算法来监控路况、司机行为

和事故，还提供了一个“撞车再现”功能，用于重建和分析事故发生环境，并与保险客户合作处理索赔。

总部位于英国的 Tractable 允许保险公司将受损图像和估计数据上传到索赔管理平台。“人工智能评论”功能会将其与数千张图片进行比较，从而相应地调整价格。

有趣的是，Tractable 还瞄准了生态系统中的其他参与者，比如汽车修理工、评估人员、供应商和租车公司。

8．防伪查询

假货越来越难被发现，网上购物比以往任何时候都更容易买到假货。为了反击，各大品牌和典当行开始尝试人工智能。

从药品到手袋再到智能手机，假货是一个影响所有零售行业的问题。

有些产品的仿冒品看起来如此逼真，以至于被归类为“超级假货”。

中国快速发展的电子商务平台拼多多在其 2018 年第三季度财报电话会议上 11 次提到“假货”，称“打击假货和有问题的商家非常困难。”

“2017 年，我们主动移除 1070 万件问题产品。我们还封锁了 4000 万个链接，这些链接引发了侵权问题。我们还与 400 多个品牌合作，共同打击假货。”

——拼多多创始人兼首席执行官黄峥

各大品牌正从两个方面与假货作斗争：

•在网络世界中，识别和删除侵犯品牌商标的在线列表，如品牌名称、标识和口号；

•在现实世界中，把奢侈品手袋等假冒产品鉴定为真定义为欺骗行为。

网上造假的规模和范围是巨大而复杂的。

电子商务巨头阿里巴巴因其网站在打击假货方面做得不够而饱受批评。据报道，阿里巴巴正在利用深度学习持续扫描其平台上的知识产权侵权行为，使用图像识别来鉴定图像中的字符，并结合语义识别，可能还会监控网站产品图像中的品牌名称或口号。

造假者使用与原始品牌列表非常相似的关键词和图片，在虚假网站上销售假货，在合法市场上销售假货，并在 Instagram 等社交媒体网站上推广假货。

当一个列表被删除时，造假者可能会用不同的关键字串重新发布相同的假冒产品。

总部位于巴塞罗那的初创公司 Red Points 正在使用机器学习技术扫描网站，寻找潜在的侵权行为，并在仿冒者使用的关键词中找到模式。其客户来自化妆品、奢侈手表、家居用品和服装行业，包括 MVMT、DOPE 和 Paul Hewitt。

在现实世界中，识别赝品更棘手，也更需要人工。

当卖家把一个二手奢侈手袋放到网上出售，或者到当铺去交易时，验证过程通常需要一个鉴定专家对包进行实物检验，包括包的制作、材料、缝制图案等。

以下是 eBay 和其他公司派出鉴定专家鉴定一个奢侈手

袋所收取的费用。

但随着“超级假货”或“3A 级假货”的兴起，用肉眼几乎无法分辨它们的区别。

建立一个假货和正品的数据库，提取它们的特征，并训练人工智能算法来分辨真假，是一个繁琐的过程。

创业公司 Entrupy 与鉴定专家合作，建立了一个假货和真货的数据库，让算法受训两年。对于稀有的古董奢侈品来说，这一过程更加困难。

Entrupy 开发了一种可以连接到智能手机上的便携式显微镜。

当用户拍摄并上传产品（手袋、手表等）的照片时，人工智能算法会分析每种产品特有的微观特征，并根据已知和真实产品的数据库对其进行验证。

数据库在不断增长，但市场上还没有一套完整的产品。Entrupty 发表的一篇论文强调了一些其他的操作假设和限制。

关键在于，与仿冒者使用的制造过程（非标准化、廉价的大规模生产）相比，使用标准或规定方法制造的物品将具有类似的视觉特征。

其次，这项技术可能不适用于纳米制造的电子芯片（处于 Entrupy 的显微镜无法检测到的尺寸变化量级）。

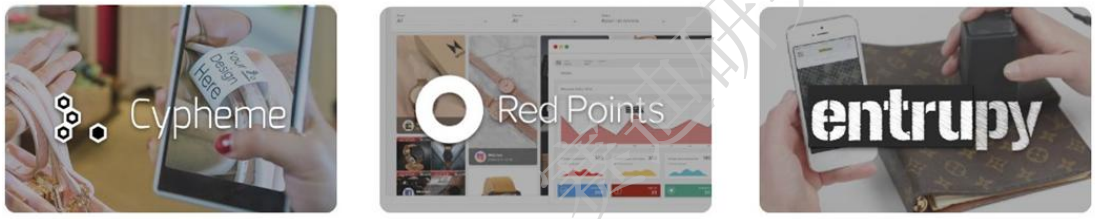

神经网络许可的随机水印用机器学习在线认证手持显微设备

Cypheme 采取了一种不同的方法，其基于水印的技术可作为贴纸贴在产品上，也可直接印在标签和包装上。

《日经亚洲评论》（Nikkei Asian Review）在对这位首席执行官的采访中详细介绍了这一技术：一滴水印生成一种随机模式。Cypheme 声称其公司拥有对这一模式的所有权，不可复制。每个独特的模式都与数据库上的特定产品相关联，使用智能手机摄像头和神经网络进行模式识别，根据数据库验证特定产品的水印模式。

这意味着 Cypheme 必须与品牌制造商直接合作，以确保产品与跟踪水印一起发货。该公司最近与欧洲领先的包装公司 AR Packaging 建立了合作关系，并与联合利华（Unilever）和雀巢（Nestle）等食品品牌合作。

虽然包装上的水印可以有效地追踪来自制造工厂和分销链上的商品，但这项技术不适用于二手采购认证。

例如，买家可以把 Cypheme 的标签从豪华手表的包装上

去掉，然后决定在经纪商店或网上转售。在这种情况下，验证真实性是不可能的，除非水印是产品本身的一部分。

对于奢侈品牌和其他高风险零售商来说，未来的解决方案可能是，在生产现场为实物商品识别或添加独特的指纹，并通过供应链对其进行跟踪。

9．免付款零售

走进一家商店，挑选你想要的东西，然后走出去，几乎“感觉”像是入店行窃。人工智能可以让真正的盗窃成为过去，让免付款零售更加普遍。

Amazon Go 取消了整个结账过程，允许购物者拿着商品走出去。

亚马逊目前还没有向其他零售商出售其“技术即服务”

（tech-as-a-service）的公开计划，而且一直对其运作、成功和痛点守口如瓶。只披露称，使用传感器、摄像机、计算机视觉和深度学习算法，并否认使用面部识别算法。

标准认知（Standard Cognition）和 AiFi 等初创企业抓住了这个机会，开始让亚马逊（Amazon）向其他零售商进军。

对于“拿了就走”的商店来说，面临的挑战是向正确的顾客收取正确的价格。

根据美国零售联合会（National Retail Federation）的数据，2017 年，由于商店行窃和文书错误等原因造成的库存损失，使美国零售商损失了约 470 亿美元。

“偷窃就是购买，”创业公司 AiFi 的联合创始人兼首席执行官 Steve Gu 在接受 AI 播客采访时谈到了“拿了就走”商店背后的技术。

到目前为止，只有 Amazon Go 在商业部署上实现了成功，但是成功的参数受到严格控制。

当你控制谁进入商店并自动付款时，入店行窃的几率就会降到最低。

亚马逊已经有了固定的 Prime 会员基础。到目前为止，所有的 Go 商店都仅限于 Prime 会员。其他 Kindle 商店等零售店，对公众开放，却仍依赖于手工结账过程。

规模较小的杂货店、便利店，甚至几家老牌超市都必须从零开始建立会员基础。

Steve Gu 在同一播客中暗示，可能会为愿意下载这款应用的人提供一个“拿了就走”的版块，为不想下载的人提供

一个单独的结账队伍。

目前还不清楚商店的基础设施将如何同时支持这两种功能。

这仍然存在销售点库存减少的问题，比如账单错误或POS机被盗。

中国的宜图科技（Yitu Technology）和东芝（Toshiba）是处理库存减少问题的两家公司，东芝拥有用于检测的智能摄像头。

防止盗窃的复杂性取决于操作的大小和规模，以及货架上产品的类型。

Amazon Go商店的面积只有1800～3000平方英尺，并使用数百个摄像头覆盖几乎每一英寸的天花板空间。相比之下，传统超市的面积可达4万平方英尺或更多。

Amazon Go除了使用摄像头进行视觉识别外，还在货架上安装了重量传感器。目前提供的商品选择有限，比如现成的和包装好的套餐。

一些需要考虑的事情是，如何利用楼面空间，尤其是在人口密集的超市，以确保相机放置于最佳位置，好用于追踪人和商品。

松散的蔬菜和其他农产品在按磅计算时可能要依靠传感器技术，但多个购物者同时从同一个纸箱中挑选商品，单靠传感器是行不通的。即使是预先包装好的或切成丁的蔬菜，价格也会因包装的不同而略有不同。

服装也特别难以被计算机视觉系统跟踪。确定衣服的尺

寸（S/M/L）和追踪容易折叠并收起来的衣服是难点所在。

STANDARD COGNITION

Standard Cognition has partnered with Japan's CPG wholesaler Paltac Corporation

AIFI

AiFi reportedly has around 20 retail clients in the pipeline

虽然初创公司 AiFi 承诺将利用现有的存储基础设施，以及传感器和摄像头的组合，但 Standard Cognition 声称将完全抛弃传感器，只依靠机器视觉。

Standard Cognition 宣布与日本最大的 CPG 批发商帕尔塔公司（Paltac Corporation）合作，在 2020 年东京奥运会之前为 3000 家日本门店提供服装。据报道，AiFi 约有 20 家零售客户正位于筹备名单中，其中包括签了合同的纽约一家大型零售商。

在短期内，问题将归结于潜在技术故障造成的部署成本和库存损失成本，以及零售商能否承担这些成本和风险。

10. 后台办公自动化

人工智能使管理工作自动化，但数据的多样性和格式使其成为一项具有挑战性的任务。

自动化“后台任务”的挑战可能是独特的，这取决于行业和应用程序。以临床试验为例，许多试验仍然依赖以纸质方式每日输入患者数据。这些记录都是数码储存的，通常情况下格式都难以搜索。而手写临床笔记对自然语言处理算法

提取信息（解释拼写错误、术语、缩写和缺少条目）提出了独特的挑战。

另一方面，自动化汽车索赔，在评估受损并探究根本原因时面临一系列不同的挑战。

但是不同的部门开始在不同程度上采用基于 ML 的工作流解决方案。

机器人流程自动化（RPA）是一个松散的术语，指的是任何后台工作都是重复的，并且可以由机器人自动完成。但是，与人工智能一样，也是一个涵盖了从数据输入到遵从性、事务处理到客户引导等一系列任务的总称。

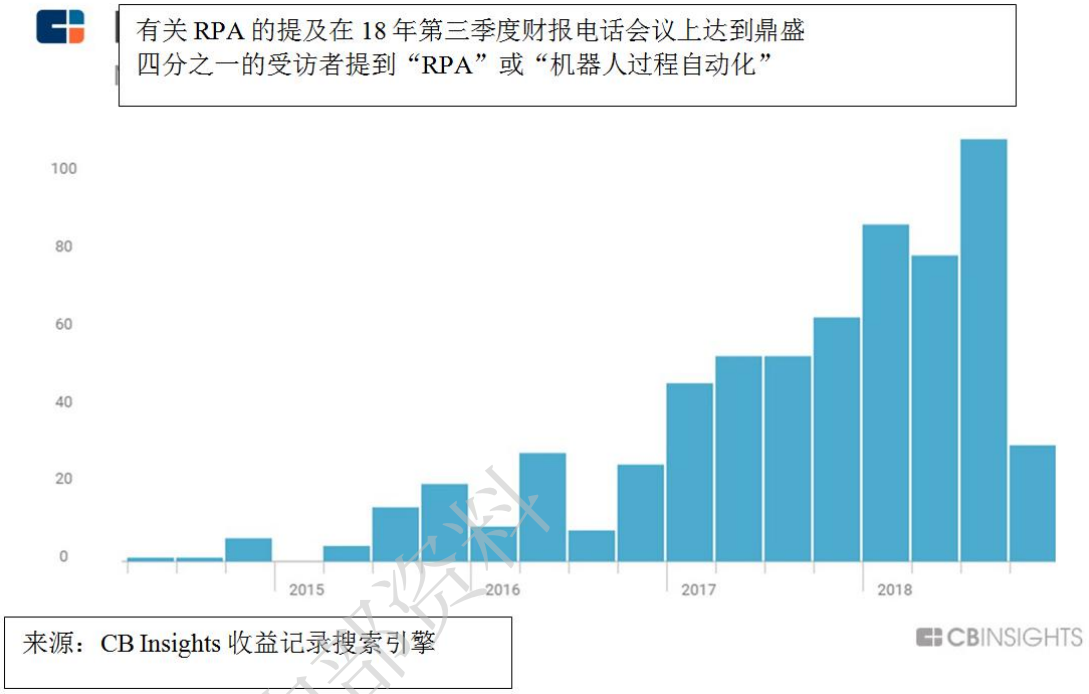

来源：CB Insights 收益记录搜索引擎

虽然不是所有的 RPA 都是基于 ML，但是很多都开始把图像识别和语言处理整合到解决方案中去。

例如，WorkFusion 自动完成后端操作，如了解客户（KYC）和反洗钱（AML）流程。

UiPath 的服务已经被全球超过 700 家企业客户使用，其中包括 DHL、NASA 和惠普，涵盖从金融到制造到零售各个行业。

Automation Anywhere 是 RPA 领域的另一个独角兽。该公司的一个案例研究突出表明，它与一家全球银行建立了合作关系，利用机器学习实现人力资源管理自动化。“IQ 机器人”从来自多个国家和多种语言的表单中提取信息，清理数据，然后自动将其输入人力资源管理系统。

尽管 RPA 的概念已经存在多年，但许多行业刚刚开始克服惯性，尝试新的技术。在其他领域，出现预测分析层之前需要进行数字化。

11．语言翻译

语言翻译的自然语言处理是一个挑战，也是一个尚未开发的市场机遇。大型科技公司正挑战这一壁垒。

基于机器的语言翻译是一个尚待开发的巨大机会，应用于跨国公司、客户支持、新闻媒体以及其他项目的后台办公自动化事宜。

最近百度宣布，将推出新的翻译耳机，类似于谷歌像素耳机，据称可以实时翻译 40 种不同的语言。

Unbabel 等一些初创公司正在使用“人在回路”机器翻译系统，目标是通过反馈回路训练算法，使其随着时间的推移变得更好。

用于翻译的 NLP 有几个挑战。例如，仅汉语的自然语言处理就很复杂，有 130 种口语和 30 种书面语。

深度学习领域的先驱研究员约舒亚•本乔发表了一篇机器翻译新架构的论文，一种使用神经网络代替传统统计方法的新方法。一年后，谷歌升级了自己的算法，用于谷歌翻译工具。

谷歌首席执行官桑达尔•皮查伊（Sundar Pichai）在2016年的财报电话会议上表示：“这一突破将帮助我们为世界各地的人们提供更准确的翻译。”

谷歌想放弃过去基于短语的机器翻译（PBMT）系统，提出了新的谷歌神经翻译系统（GNTS）。

虽然神经机器翻译方面发表了许多论文，但也存在一些局限性，比如训练这些模型所需的时间和计算资源，以及翻译不出罕见词汇。

谷歌提出了解决这些问题的改进建议，并测试了其在英到中、中到英、西班牙语到英语等方面的算法。

Language pair	***Input sentence:***	***Translation (PBMT):***	***Translation (GNMT):***	***Translation (human):***
Chinese->English	2015年到2016年，亚太地区的信用卡交易总额将增加1.7万亿美元，增幅最大，其次是北美地区，增幅将达到1870亿美元。	2015 to 2016 , the total amount of credit card transactions in Asia Pacific will increase by $ 1.7 trillion , the largest increase , followed by North America , growth will reach $ 187 billion .	Total credit card transactions in the Asia-Pacific region will increase by $ 1.7 trillion in 2015-2016, the largest increase, followed by North America with $ 187 billion.	Total credit card transactions in the Asia-Pacific region will increase by $ 1.7 trillion in 2015-2016, the largest increase followed by North America with $ 187 billion.
Chinese->English	100年前，预测引力波的爱因斯坦或许都无法想象人类可以直接观测到引力波。	100 years ago, the prediction of Einstein's gravitational waves probably can not imagine humans can directly observe gravitational waves.	100 years ago, Einstein predicted gravitational waves may not be able to imagine humans can directly observe the gravitational waves.	100 years ago, Einstein who predicted gravitational waves may not be able to imagine that humans can directly observe the gravitational waves.

Source: Google AI Blog

关于这个主题已经发表了几篇研究论文。但最近的突破

来自 Facebook。

根据这篇论文，“对多语言 NLP 的研究大多集中在高资源语言，如汉语、阿拉伯语或主要的欧洲语言，通常仅限于少数几种（通常只有两种）。

相比之下，我们学习 93 种不同语言的联合表示方法，这些语言包括资源缺乏的语言和少数民族语言。”

		Details			Training corpus size	Tatoeba Error [%]		Tatoeba test set size
ISO3	ISO2	Name	Family	Script		en → xx	xx → en	
hye	hy	Armenian	Armenian	Armenian	6k	59.97	67.79	742
bel	be	Belarusian	Slavic	Cyrillic	5k	31.20	36.50	1000
mya	my	Burmese	Sino-Tibetan	Burmese	2k	n/a	n/a	–
dtp		Central/Kadazan Dusun	Malayo-Polynesian	Latin	1k	92.10	93.50	1000
khm	km	Central Khmer	Khmer	Khmer	625	77.01	81.72	722
cbk		Chavacano	Creole, Romance	Latin	1k	24.20	21.70	1000
kzj		Coastal Kadazan	Malayo-Polynesian	Latin	560	91.60	94.10	1000
cor	kw	Cornish	Celtic	Latin	2k	91.90	93.20	1000
mhr		Eastern Mari	Uralic	Cyrillic	1k	87.70	91.50	1000
ido	io	Ido	constructed	Latin	3k	17.40	15.20	1000
ina	ia	Interlingua	constructed	Latin	9k	5.40	4.10	1000
ile	ie	Interlingue	constructed	Latin	3k	14.70	12.80	1000
gle	ga	Irish	Irish	Latin	732	93.80	95.80	1000
kaz	kk	Kazakh	Turkic	Cyrillic	4k	80.17	82.61	575
lfn		Lingua Franca Nova	constructed	Latin	2k	35.90	35.10	1000
oci	oc	Occitan (post 1500)	Romance	Latin	3k	39.20	38.40	1000
wuu		Wu Chinese	Chinese	Chinese	2k	25.80	25.20	1000
yue		Yue Chinese	Chinese	Chinese	4k	37.00	38.90	1000

Table 2: List of the 18 very low-resource languages included during training of the proposed model, along with their language family, writing system, the resulting similarity error rate on the Tatoeba test set, and the number of sentences in it. Dashes denote language pairs excluded for containing less than 100 test sentences.

Source: https://arxiv.org/pdf/1812.10464.pdf

随着大型科技公司继续投入资源改善翻译框架，其效率和语言能力将会提高，各行业的应用也会增加。

12．合成训练数据

要训练人工智能算法，访问大型、带标签的数据集是必要的。真实的假数据可能会解决这个瓶颈。

人工智能算法的好坏取决于输入的数据，而为不同的应用程序访问和标记这些数据是需要投入大量时间和资本的。甚至来说，访问这类真实世界的数据可能不可行。

以自动驾驶汽车为例。训练自动驾驶汽车应对不太常见的危险情况，比如刺眼的太阳下或行人从停着的汽车后面跳出来，使用真实数据很难。

这就是合成数据集的用武之地。

2018 年 3 月，英伟达推出了一款基于云计算的自动驾驶汽车仿真软件，名为 DRIVE Constellation。自动驾驶汽车在上路前可以在虚拟现实模拟中行驶数十亿英里，且不发生事故——这是一项旨在创造“一种更安全、更可扩展的方法，让自动驾驶汽车上路。”

想象一下自动驾驶汽车在雷暴中行驶。英伟达的解决方案模拟了汽车中的数据传感器（比如摄像头或激光雷达）在这些条件下会产生什么变化。合成传感器数据被输入计算机，计算机就像在实际道路上行驶一样做出决策，并向虚拟车辆发回命令。

一个有趣的新兴趋势是使用人工智能本身，来帮助生成更“真实”的合成图像来训练人工智能。

例如，英伟达使用生成性对抗网络（GANs）来创建带有脑肿瘤的假 MRI 图像。

“总的来说，这些结果为机器学习在医学成像领域面临的两大挑战提供了一个潜在的解决方案，这两大挑战分别是：病理发现的发生率较低，以及共享患者数据方面的限制。”

——英伟达研究论文

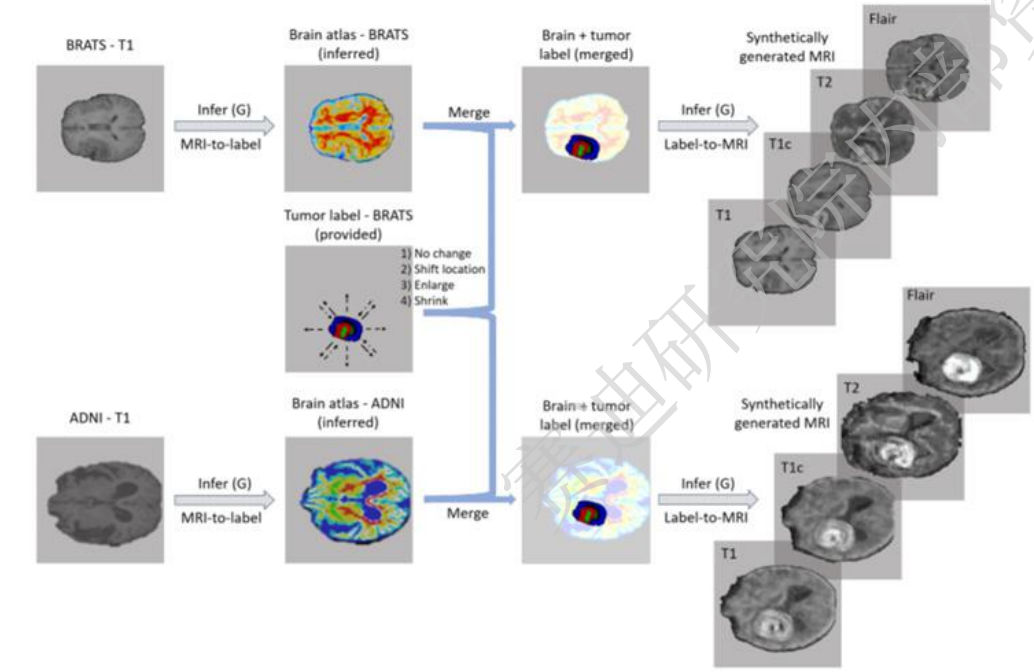

Source: https://arxiv.org/pdf/1807.10225.pdf

GANs 被用来“扩充”真实世界的数据，这意味着人工智能可以通过混合真实世界和模拟数据进行训练，从而拥有更大、更多样化的数据集。

机器人技术是另一个可以从高保真度合成数据中获益良多的领域。

考虑一个简单的任务，教机器人抓取东西。2016 年，谷歌的研究人员使用了 14 条机器手臂，任务是学习如何抓住不同的物体。来自 14 个机器手臂的失败和成功尝试的数据被用来训练一个神经网络，帮助机器人“分享经验”并预测抓取的结果。

研究团队称，总共需要 80 万次抓握尝试，“相当于机器人 3000 小时的练习”，才能“看到智能反应行为的开始出现”。

但是模拟——让数百个虚拟机器人在虚拟环境中练习——可以极大地简化这个过程。

其中一个挑战是创建逼真的对象（比如使苹果或铅笔的模拟看起来尽可能接近真实的对象）。2017 年，谷歌的研究人员利用生成式对抗网络（GANs）来实现这一点，大大减少了训练机器人所需的真实世界数据量。

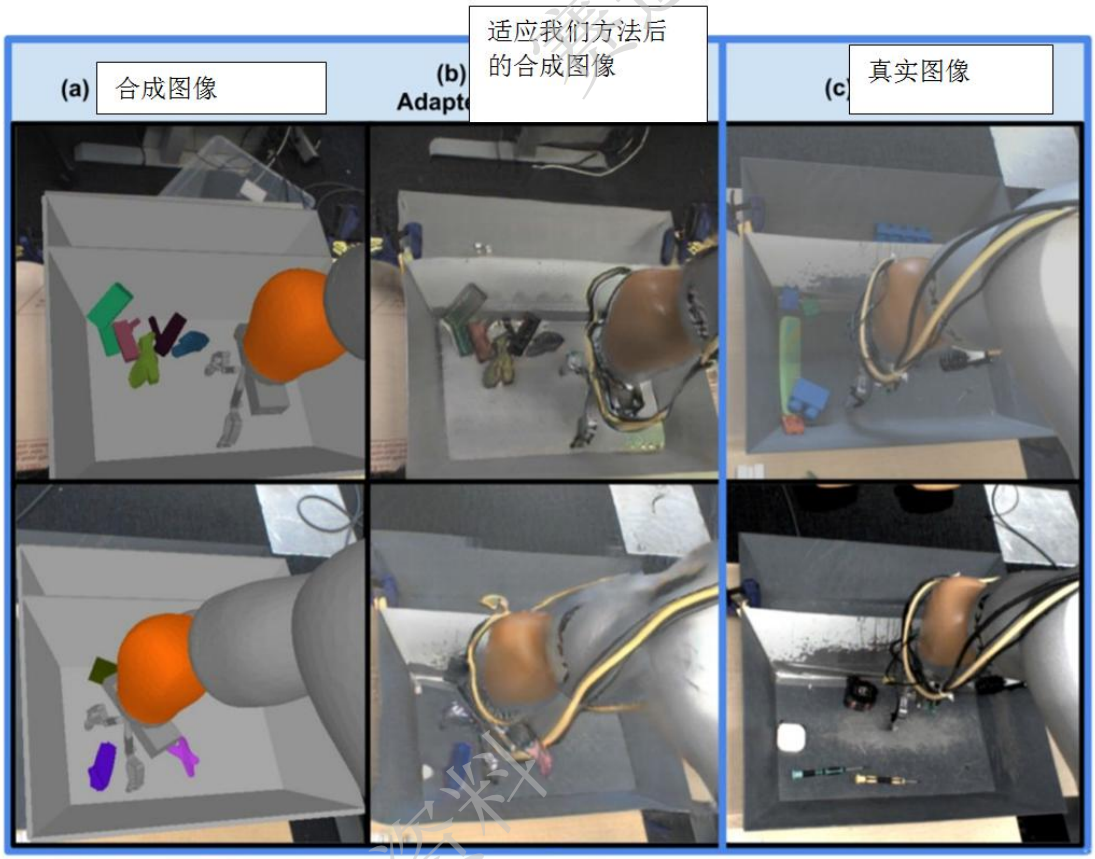

Source: https://arxiv.org/pdf/1709.07857.pdf

像 AI.Reverie 这样的早期创业公司正在开发仿真平台，为各种行业和场景生成数据集。

随着技术规模的扩大，以及合成数据能更准确地模拟现实场景，合成训练数据将成为无法访问大型数据集的小型公司的催化剂。

威胁性

1．强化学习

从训练算法到在棋类游戏中击败世界冠军，再到教授人工智能杂技算法，研究人员正在用强化学习来突破边界。但是目前对大量数据集的需求限制了实际应用。

当谷歌Deep Mind的Alpha Go在复杂而具有战略意义的中国围棋比赛中击败一名世界冠军时，强化学习得到了媒体的关注。

强化学习的要点是：你需要采取什么行动才能达到你的目标并获得最大的回报?

因为这种方法，强化学习在游戏和机器人仿真中特别流行。

Deep Mind 的 Alpha Go 最初是使用监督学习（使用来自其他人类玩家的数据来训练算法）和强化学习（人工智能对抗自己）进行训练。

Deep Mind 后来发布了 Alpha Go Zero，并声称已经实现了超人的性能，完全是基于强化学习训练出来的（只给出一组规则来对抗自己）。

最近，加州大学伯克利分校的研究人员利用计算机视觉和强化学习，从 YouTube 视频中教授杂技算法技巧。计算机模拟的角色可以复制视频中的动作，而不需要手动标注姿势。

Source: https://bair.berkeley.edu/blog/2018/10/09/sfv/

通过强化学习，模拟人物可以将自己的技能应用到新的环境中。例如，如果 YouTube 视频中的一名男子在平地上做了后空翻，这个模拟的角色可以调整技能，在不平坦的地形上做后空翻。

尽管有了这些快速的进步，强化学习的应用还没有开始，因为与监督学习相比，需要大量的数据，而监督学习是当今最流行的人工智能范式。

“当你想到如今创造的经济价值时，清单（不同的学习方法）上的数字会迅速下降……强化学习是一种技术，在这种技术中，公关的兴奋程度与当前的实际部署极不相称。”

——ANDREW NG， EMTECH 2017 简报

但是对 RL 应用的研究正在增加。对美国专利申请标题和摘要关键词的搜索显示，过去两年，专利申请活动有所增加。

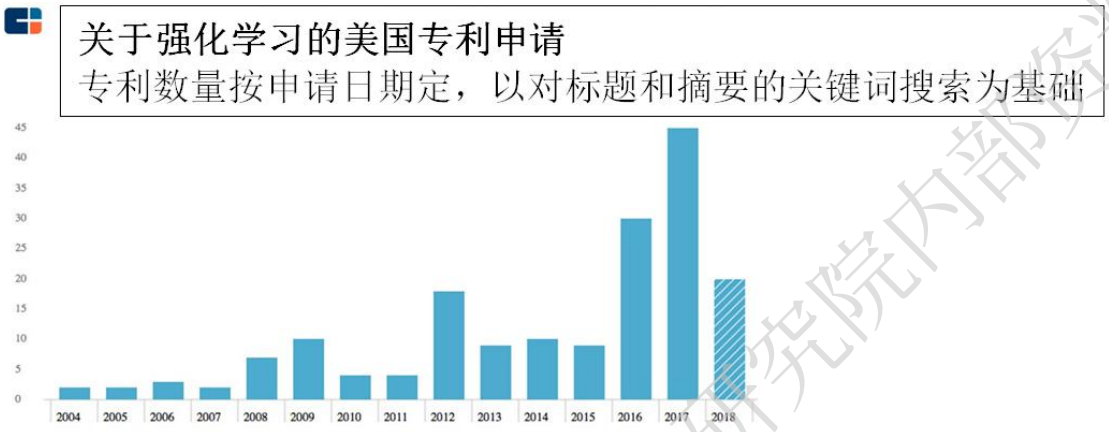

注释：由于专利申请和发表有延迟，柱状图可能呈数量减少趋势
来源：CB Insights 专利分析

最热门的申请者包括谷歌、IBM、Alphaics（一家人工智能初创企业）、Mobileye（被英特尔收购）、微软、Adobe 和发那科。

在财报电话会议中，百度积极讨论强化学习，在 2018 年第一季度电话会议中 7 次提到了强化学习。

“一季度的亮点之一是，我们首次部署了一个强大的基于强化学习的基础设施，可以显著提高我们更好地匹配用户广告的能力，并提高点击率和转化率。”

——百度在 2018 年第一季度财报电话会议

2. 网络优化

从促进频谱共享到监控资产，再到天线优化设计的出现，人工智能开始改变通信。

电信网络优化是一组技术，用于改善延迟、带宽、设计或体系结构——以一种有利的方式增强数据流。

对于通信服务提供商，优化直接转化为更好的客户体验。

除了带宽限制，电信领域最大的挑战之一是网络延迟。手机上的 AR/VR 等应用程序只能以极低的延迟时间优化功

能。

最近，苹果获得了一项专利，可以利用机器学习来形成“预期网络”，预测智能手机等支持无线的设备未来可能会采取何种行动，并提前下载数据包，以减少延迟。

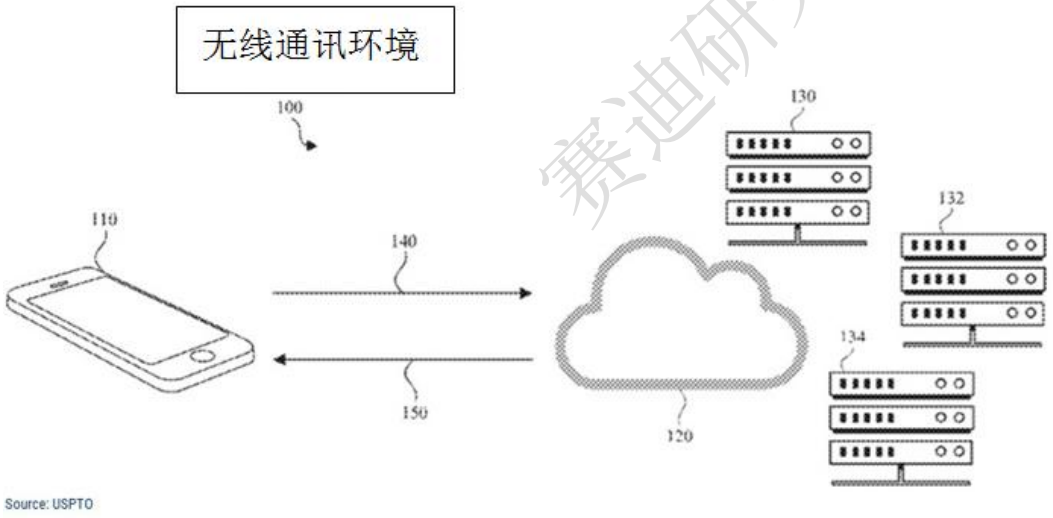

Source: USPTO

机器学习的另一个新兴应用是频谱共享。

政府在拍卖中向 Verizon 这样的公司发放了特定频率的电磁频谱许可证。

联邦通信委员会（FCC）规定 3.5 到 3.7GHz 的频谱将在不同用户之间共享。

这意味着运营商可以根据可用性动态访问共享频率。这将允许他们根据网络需求上下扩展带宽，还将为规模较小的商业用户提供频谱接入，而这些用户没有获得自己专用频谱的授权。

3.5GHz 频段的部分频段被美国海军和其他联邦机构使用。他们被赋予了第一层访问权限，如果他们未使用这一频谱，那么就属于第 2 层和第 3 层用户。

像 Federated Wireless 这样的公司提供安全的频谱访问（SAS），以便在不同层次的用户之间动态分配频谱，并确保不干扰联邦信号——利用机器学习来做到这一点。

2018 年，Federated Wireless 获得了一项专利，可以使用 ML 将无线电信号分为不同的类别，比如联邦信号、噪声信号和未知信号，这样做的同时模糊了联邦信号的特征（因此黑客永远无法访问军事/国防信号的特定特征或弱点）。

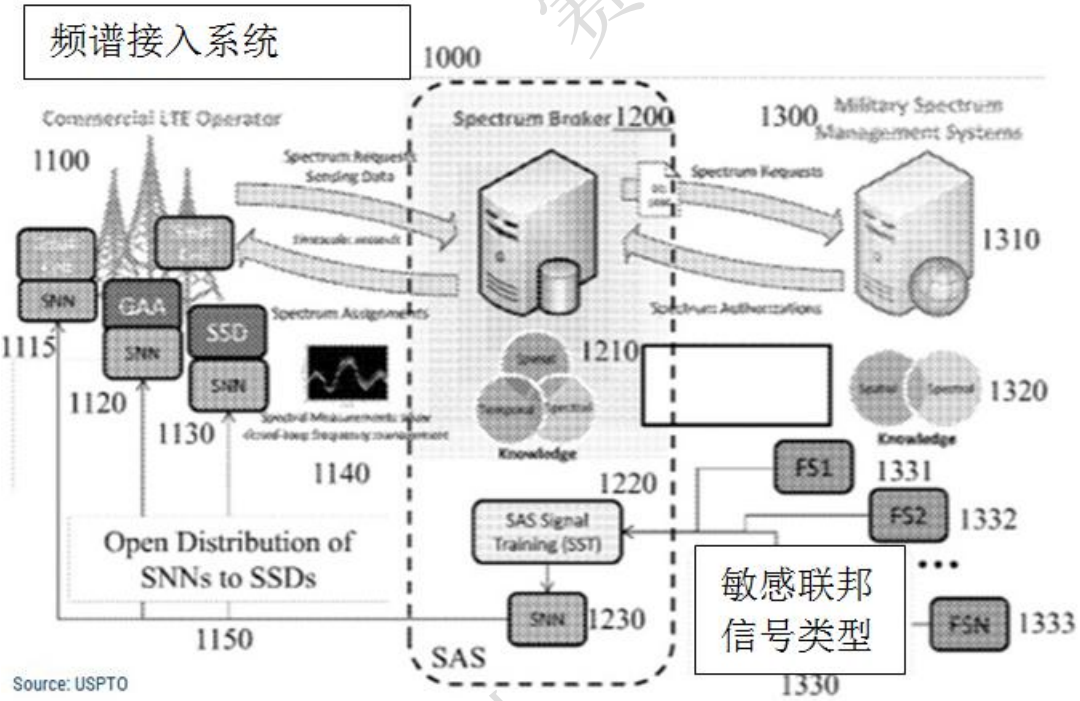

DARPA 希望最终摆脱 SAS 的角色，将频谱共享变成一个基于 ML 的自动化系统。

为此，该公司在 2016 年推出了频谱协作挑战赛。竞赛参与者必须使用 ML 来为无线网络提出独特的方法，以“自动协作，动态地确定如何每时每刻使用射频（RF）频谱”。

DARPA 还在 2017 年推出了一个射频机器学习系统（RFMLS）项目。与上面的联邦无线专利类似，DARPA 希

望使用 ML 来区分不同类型的信号，特别是识别意图侵入终端设备（如物联网设备）的恶意信号。

电信运营商也准备将基于人工智能的解决方案整合到下一代无线技术（即 5G）中。

三星收购了基于人工智能的网络和服务分析初创公司 Zhilabs，为 5G 时代做准备。

三星在一份新闻稿中表示，人工智能软件将被用于“分析用户流量，对正在使用的应用程序进行分类，并提高整体服务质量。”

高通将人工智能边缘计算视为其 5G 计划的一个关键组成部分（边缘计算减少了带宽限制和与云的频繁通信——这是 5G 的一个主要重点领域）。

早期的研究论文也在探索，使用神经网络来提出电信网络上最优的设计天线。

3．自动驾驶汽车

尽管自动驾驶汽车存在巨大的市场机遇，但完全自动的时间表仍不明朗。

许多大型科技公司和初创企业正在自动驾驶汽车领域展开激烈竞争。

谷歌在汽车领域已经负有盛名，其自动驾驶项目 Waymo 首次部署商用自动驾驶车队。

投资者对开发全自动驾驶汽车的公司仍充满信心，他们向通用汽车的巡航自动化系统投入了数亿美元（2018 年 10 月本田投资 7.5 亿美元，软银投资 9 亿美元），Zoox 投资 5

亿美元（2018 年 7 月 5 日）。这里的其他初创公司包括 Drive、小马和 Nuro。

中国在这方面尤其加大了力度。中国科技部 2018 年宣布，中国第一批开放的人工智能平台将大多依赖百度自动驾驶。

2017 年 4 月，百度宣布了一个独一无二的自主驾驶解决方案开放平台——阿波罗，吸引了全球各地的合作伙伴。

与其他开源平台一样，这一想法是通过向生态系统中的其他参与者开放人工智能和自动驾驶研究，从而加速这方面的研究。让每个人都可以使用源代码，允许公司在现有研究的基础上进行构建，而不是从零开始。

百度建立全球自动驾驶汽车生态系统
为百度阿波罗计划选择国际伙伴

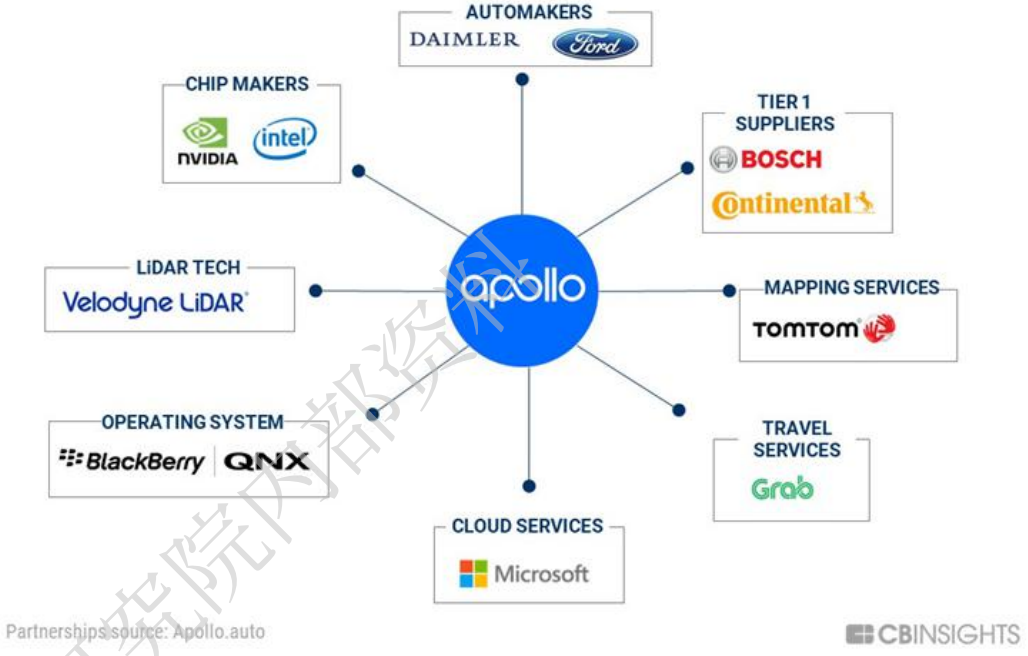

阿里巴巴最近也对其自动驾驶汽车进行了试驾。但有趣的是，就在一年前，阿里巴巴对自动驾驶汽车的长期商业机

会表示怀疑，在财报电话会议中提到“没有人找到了长期的经济模式，但人们正在研究，因为在建立自动驾驶汽车过程中有一些非常有趣的人工智能技术”。

尽管人们对这项技术的未来心存疑虑，但汽车制造商仍在全速前进。预计到 2025 年，该市场将达到 800 亿美元左右。

一些应用程序可能会更早地采用全自动驾驶汽车，比如物流和物流配送。

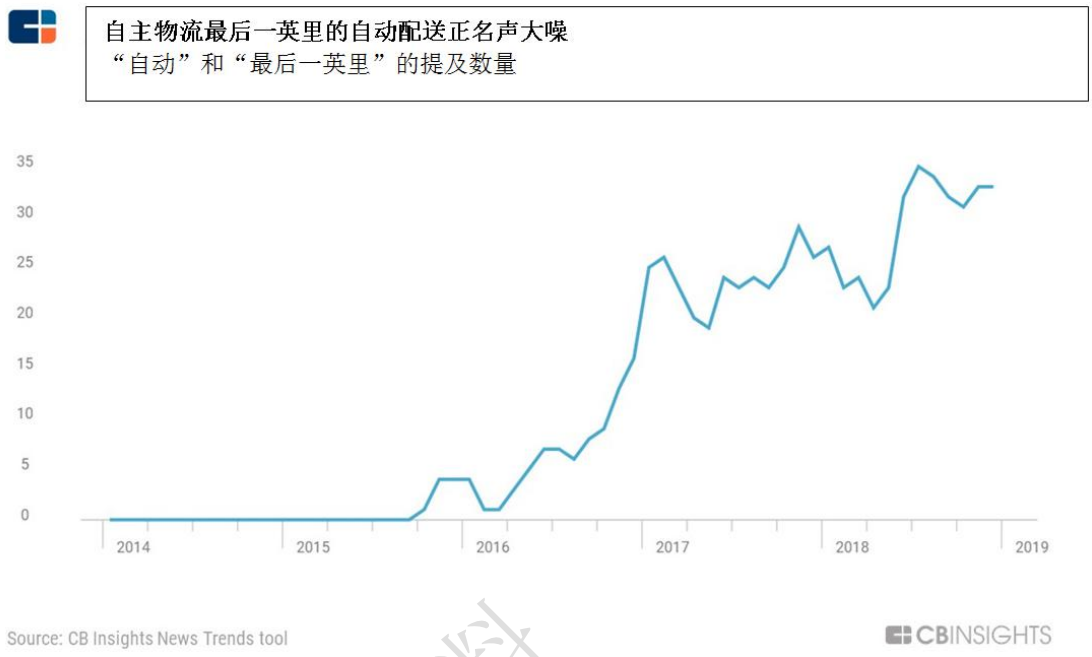

自主物流——尤其是最后一英里的自动配送——是零售商和物流公司的首要任务，而且可能是我们看到完全自主的第一个领域。自动驾驶汽车可以帮助解决在最后一英里交付货物这一昂贵而艰巨的挑战，这可能会使一件商品的总交付成本增加近 1/3。

对自动驾驶车辆部署实施宽松管理的亚利桑那州正在

成为试验基地。2018 年 6 月，机器人初创公司 Nuro 与美国最大的实体杂货商之一 Kroger 合作配送食品杂货。Nuro 的设计初衷是在社区道路上行驶，而不仅仅是像其他已经开发出来的送货机器人和车辆原型那样在人行道上行驶。

在餐厅领域，达美乐（Domino’s）和必胜客（Pizza Hut）等披萨公司一直在测试自动驾驶汽车。福特正在迈阿密试行自动配送披萨、杂货和其他商品。2018 年初，OEM 与包括达美乐在内的 70 多家企业合作。

4．作物监测

三种类型的作物监测正在农业领域起飞：地面、空中和地理空间。

到 2021 年，精准农业无人机市场预计将达到 29 亿美元。

无人机可以为农民绘制农田地图，利用热成像技术监测水分含量，识别受虫害影响的作物并喷洒农药。

初创企业正专注于为第三方无人机捕获的数据再增加一层分析功能。

例如，塔拉尼斯公司就使用第三方塞斯纳飞机来完成这一任务。塔拉尼斯 2018 年还收购了农业科技 AI 初创公司 Mavrx Imaging，该公司当时正在开发超高分辨率成像技术，用于侦察和监控油田。

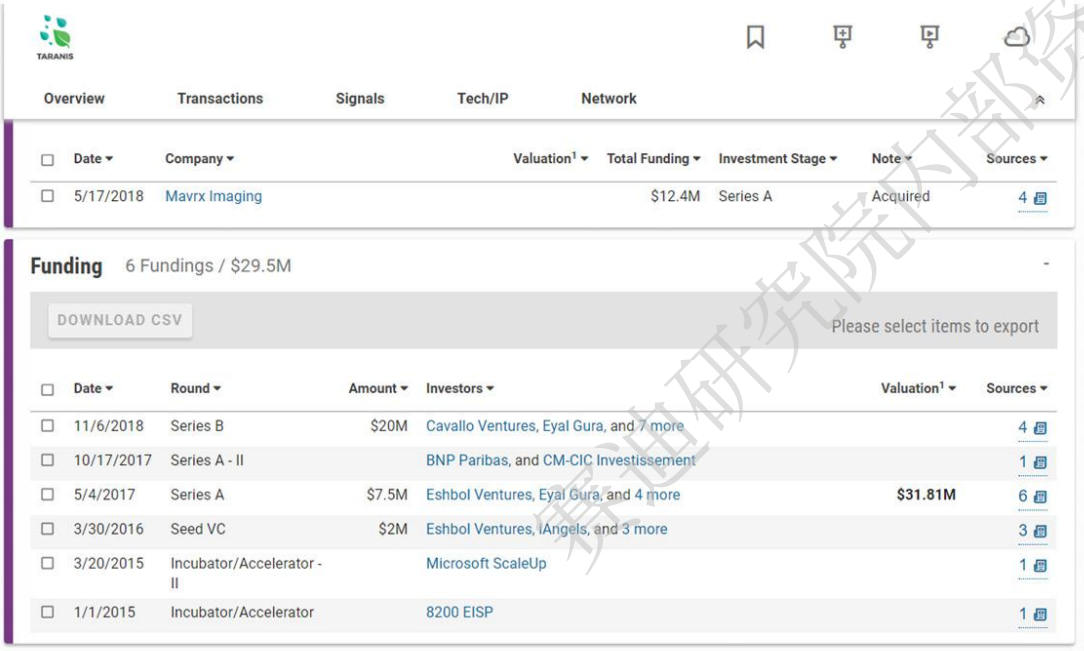

Source: CB Insights

CBINSIGHTS

塔拉尼斯使用人工智能将田地的图像拼接在一起，还能识别出农作物可能存在的问题。农业设备制造商约翰迪尔（John Deere）正与其他几家公司合作，开发了潜在的解决方案。

约翰迪尔一直在用人工智能重塑自己，已收购了蓝河科技——一家利用计算机视觉的农业设备公司——投资 3 亿美元以上。除此之外，蓝河科技还在致力于“智能除草”和“see-and-spray”解决方案。

这种单独的作物监测可能成为农业农药工业的主要干扰因素。如果地面上的农业设备能更智能地利用计算机视觉，并根据需要只喷洒个别作物，这将减少对非选择性除草剂的需求，而非选择性除草剂会杀死附近的一切作物。精确喷洒还意味着减少除草剂和杀虫剂的使用量。

在该领域之外，利用计算机视觉分析卫星图像，提供了对农业实践的宏观理解。

地理空间数据可以提供关于全球作物分布格局和气候变化对农业的影响的信息。

嘉吉投资了笛卡尔实验室（Descartes Labs），该实验室利用卫星数据开发了大豆和玉米等作物的预测模型。这个应用程序的计算机视觉也引起了大宗商品交易商和政府机构的兴趣。DARPA 正与笛卡尔合作预测粮食安全。

暂时性

1. 网络威胁追踪

仅仅对网络攻击做出反应是不够的。在网络安全领域，主动利用机器学习“追踪”威胁的势头正在增强。

计算能力和算法的进步正使以前理论上的黑客攻击变成真正的安全问题。

根据全球公共数据泄露数据库 Breach Level Index， 18 年上半年全球共有 45 亿份数据记录被泄露（供参考，2017 年全年的数据为 26 亿份）。

Rank	Organization Breached	Records Breached	Date of Breach	Type of Breach	Source of Breach	Location	Industry	Risk Score
1	Facebook	2,200,000,000	04/04/18	Identity Theft	Malicious Outsider	United States	Social Media	10.0
2	Exactis	340,000,000	06/01/18	Identity Theft	Accidental Loss	United States	Other	9.1
3	Under Armour	150,000,000	02/01/18	Account Access	Malicious Outsider	United States	Retail	9.1
4	Twitter	336,000,000	05/03/18	Financial Access	Accidental Loss	United States	Social Media	9.0
5	Firebase (Google)	100,000,000	06/20/18	Identity Theft	Accidental Loss	United States	Technology	8.6

Source: Breach Level Index top H1'18 breaches

与人工智能的其他工业应用不同，网络防御是一种黑客和安全人员之间的猫鼠游戏，两者都利用机器学习的进步来提升自己的游戏水平，并保持领先。

威胁追踪，顾名思义，是一种主动发现恶意活动的行为，而不仅仅是对警报或入侵事件做出反应。

追踪从一个关于网络潜在弱点的假设开始，然后用手工和自动化工具，在一个连续的、迭代的过程中测试这个假设。网络安全领域庞大的数据量使得机器学习成为这个过程中

不可分割的一部分。

在 Linkedin 上快速搜索 “威胁猎手”，就会看到微软、雷神、威瑞森、博思艾伦和道琼斯等公司在美国发布的 70 多个职位列表。

这反映了不同商业类型对威胁猎手的需求不断增长，但也表明了游戏本身仍然是利基市场。

“SANS 2018 年威胁追踪调查的结果显示，对于许多组织来说，追踪仍然是一个新事物，从过程和组织的角度来看，定义也很模糊……这项对 600 名受访者的调查显示，大多数正在追踪的组织大多是大型企业或那些在过去被当成靶子瞄准的企业。”

——IBM 发起的 SANS 2018 调查

正如 SANS 2018 年的调查所显示的那样，大型企业的风险更高，因为它们的差异化因素是对宝贵的数据的获得权。

例如，亚马逊就面临着来自 AWS 客户越来越大的保护云的压力。错误配置的 AWS 服务器导致 Verizon、WWE、道琼斯和埃森哲等客户的数据泄露。

亚马逊收购了威胁追踪初创公司 Sqrrl，以开发一种新产品，用于在 AWS 客户的账户上搜索黑客。

另一家专注于威胁追踪的人工智能初创企业 Cylance2018 年被黑莓（Blackberry）收购。

网络越分散，就越容易受到攻击。

威胁追踪可能会得到进一步的发展，但它自己也面临一系列挑战，比如应对不断瞬息万变的动态环境和减少误报。

2．对话式人工智能

对许多企业而言，聊天机器人已成为人工智能的同义词——但这一承诺并没有跟上现实。

最近，谷歌的对话式人工智能功能 Duplex 陷入了困境。

Duplex 可以代表用户打电话和做出预订，但像真人一样交流（带有“嗯”和停顿），引发了人们对 Duplex 是否需要在与真人交谈时将自己定位为对话代理的伦理担忧。

谷歌在新款手机 Pixel 3 上增加了双工功能，把 Pixel 3 变成了一个人工智能发电站，包括一个“屏幕呼叫”选项，允许谷歌助手对垃圾邮件呼叫者进行筛选。

自 2014 年以来，谷歌一直在为两个会话代理之间的交互申请专利。最近的一次申请是在 2018 年 4 月提交的，名为“对话代理响应由情绪决定”。

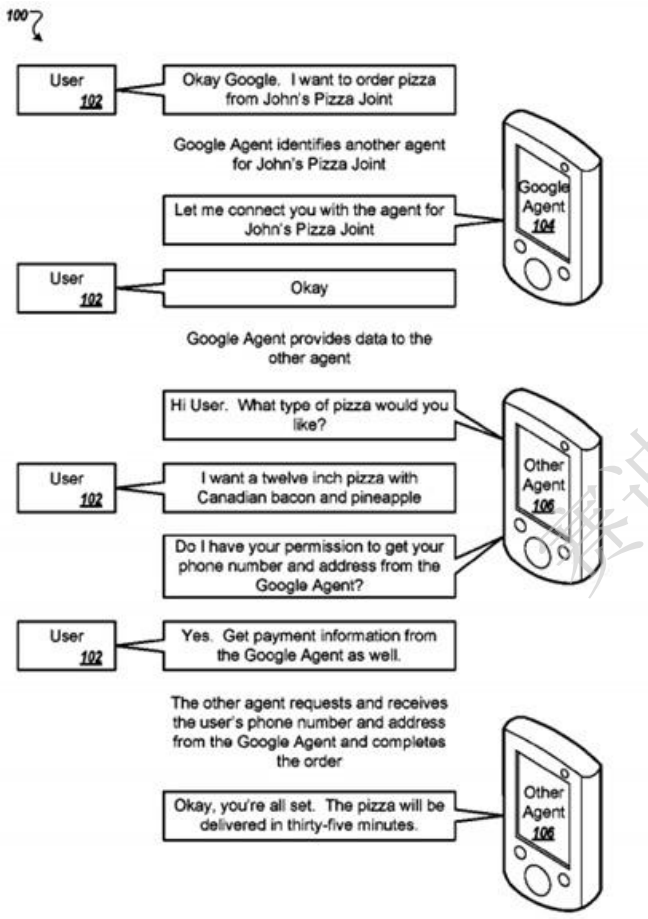

Source: USPTO

尽管 FAMGA 和中国的大型科技公司（百度、阿里巴巴和腾讯）高度关注这一领域，对话代理（语音和文本）在一些应用中比其他应用更可行。

聊天机器人最广泛的应用之一是客户服务。机器人构成与用户交互的第一层（注意：并非所有机器人都使用自然语言处理），并根据复杂程度将查询传递给人工。

对于医疗和保险等应用程序来说，这仍然是一个挑战，因为在这些应用程序中，分类（评估情况的紧急程度）是复杂的。

同样，在没有视觉提示的情况下，仅通过语音对话购物

也是一项挑战。

尽管从丝芙兰（Sephora）、雀巢（Nestle）到凯捷（Capgemini）等分析师和 CPG 品牌都表示，语音购物将成为零售业的下一件大事，但它并没有取得成功。除了重新订购特定商品外，语音购物未能提供推动在线商务发展的关键客户体验。

精神卫生保健是聊天机器人似乎具有潜在破坏力的另一个领域。

高昂的心理健康治疗成本和 24 小时可用性的吸引力，正在催生一个基于人工智能的心理健康机器人新时代。早期的初创公司专注于使用认知行为疗法——改变消极的想法和行为——作为市场上许多情绪跟踪和数字日记健康应用的对话扩展。

但心理健康是一个范围。症状上具多变性，分析时带主观性，需要高水平的情感认知和人与人之间的互动。

这使得像精神卫生保健这样的领域——尽管有成本和可访问性的优势——对算法来说尤其困难。

3．药物开发

随着人工智能生物技术初创企业的出现，传统制药公司正期待着人工智能 SaaS 初创企业为漫长的药物研发周期提供创新的解决方案。

2018 年 5 月，辉瑞与 XtalPi 达成战略合作伙伴关系。XtalPi 是一家由腾讯和谷歌等科技巨头支持的人工智能初创公司。两家企业将预测小分子的药物特性，并开发“基于计

算的理性药物设计”。

但辉瑞并非唯一的市场参与者。

诺华（Novartis）、赛诺菲（Sanofi）、葛兰素史克（GlaxoSmithKline）、安进（Amgen）和默克（Merck）等顶级制药公司近几个月都宣布与人工智能初创企业建立合作关系，以发现从肿瘤学到心脏病等一系列疾病的候选新药。

“我们仍处于早期阶段的最大机会是，利用深度学习和人工智能来识别全新的适应症，全新的药物。”

——BRUNO STRIGINI，诺华肿瘤公司前首席执行官

对该领域的兴趣正在推动人工智能药物研发初创企业的股权交易数量：截至 2018 年第二季度，共有 20 宗，相当于 2017 年全年的交易数量。

大型医药公司兴趣陡增，促进人工智能药物研发交易数量增加
2014~2018 年间股权交易

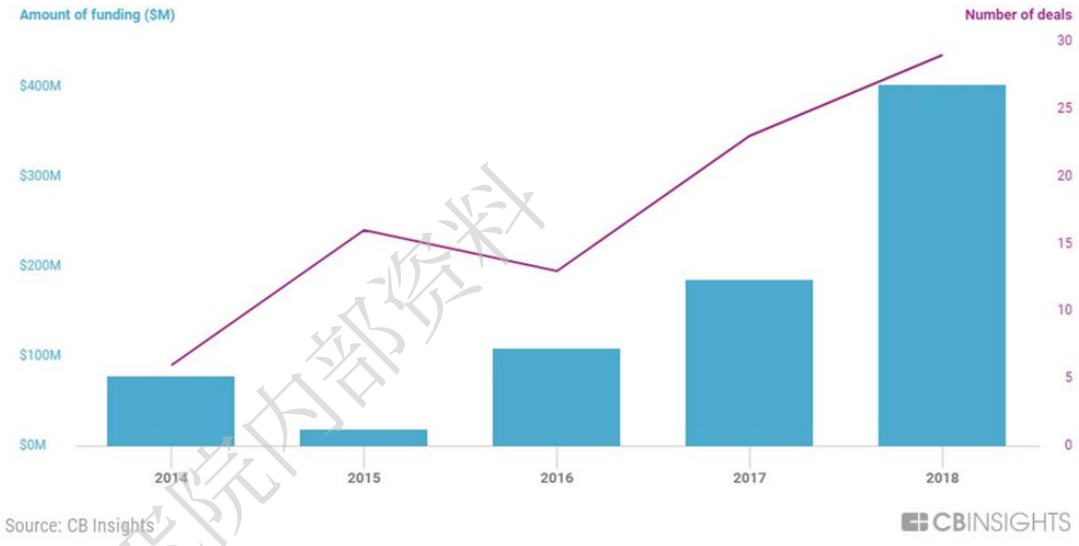

像递归制药（Recursion Pharmaceuticals）等的生物技术人工智能公司同时投资于人工智能和药物研发，而传统制药

公司则与人工智能 SaaS 初创公司合作。

尽管这些初创公司中有许多仍处于早期融资阶段，他们已经拥有一批制药客户。

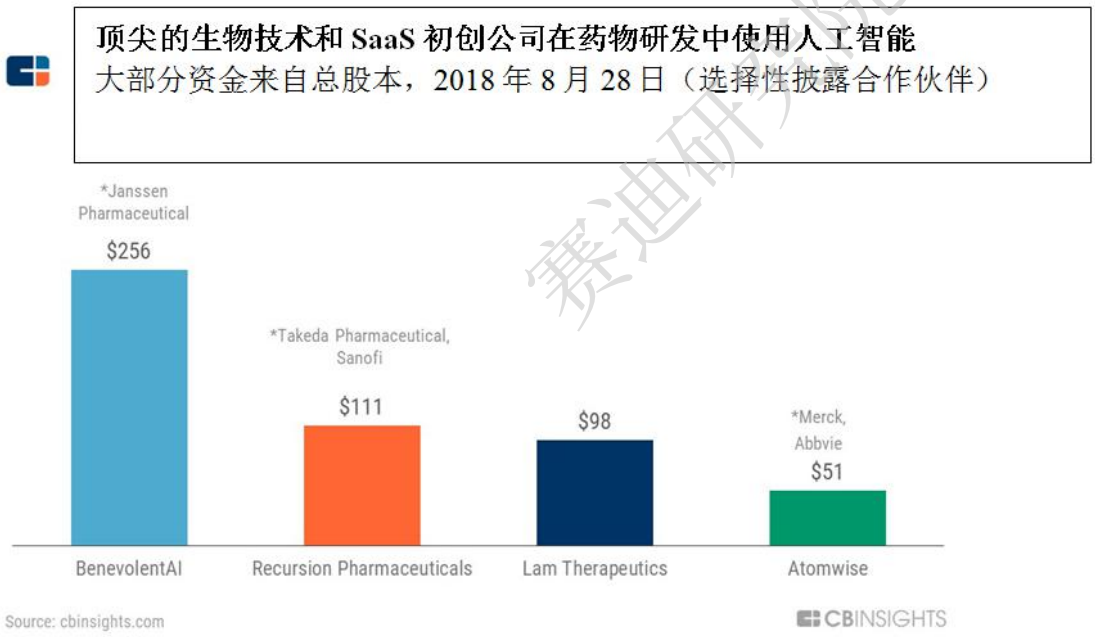

在药物配方阶段，很少有可衡量的成功指标，但制药公司正将数百万美元押在人工智能算法上，以发现新的治疗候选药物，并改变旷日持久的药物发现过程。

这些数据从何而来?

CB Insights 平台的基础数据包含在该报告中

点击这里免费注册

CBINSIGHTS

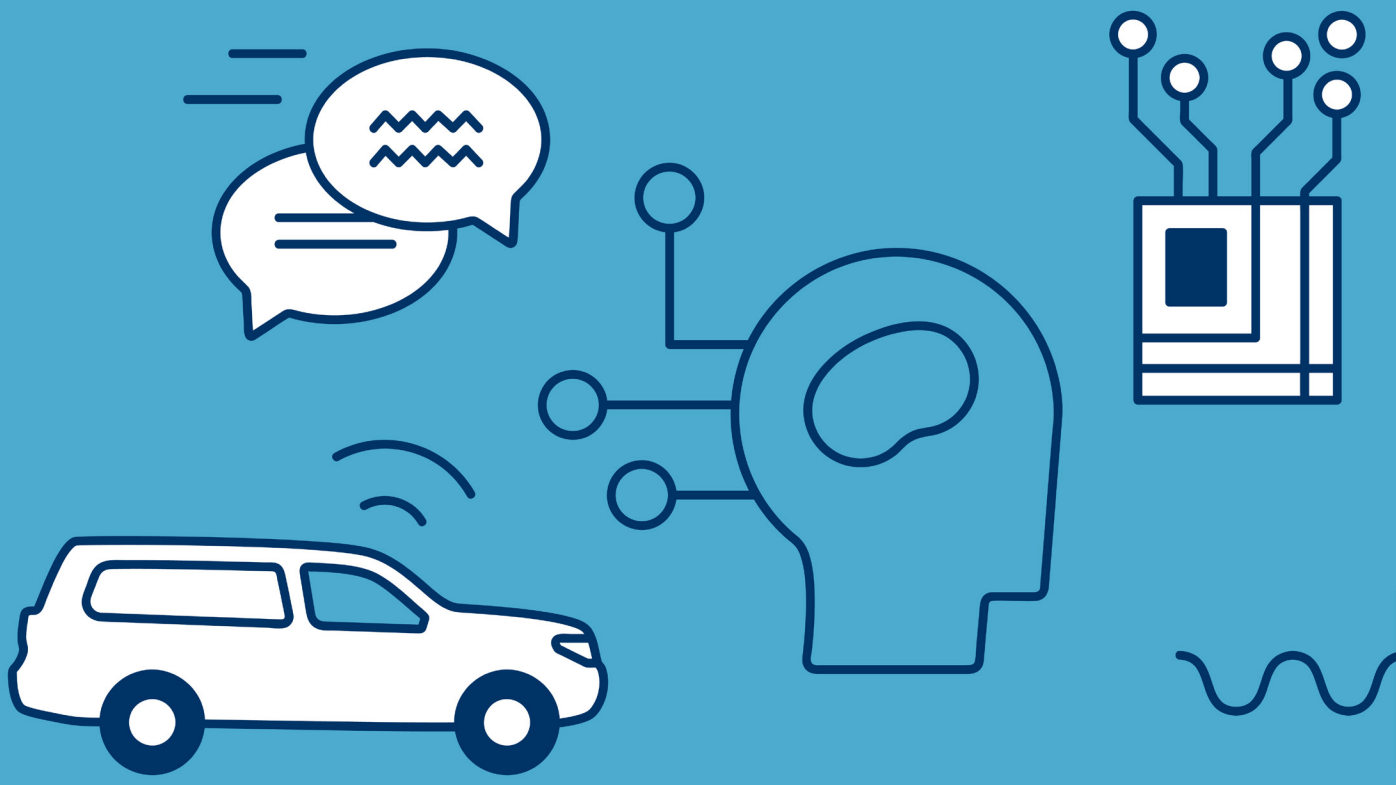

WHAT'S NEXT IN AI?

Artificial Intelligence Trends

2019

Table of Contents

NExTT FRAMEWORK

Artificial Intelligence Trends in 2019

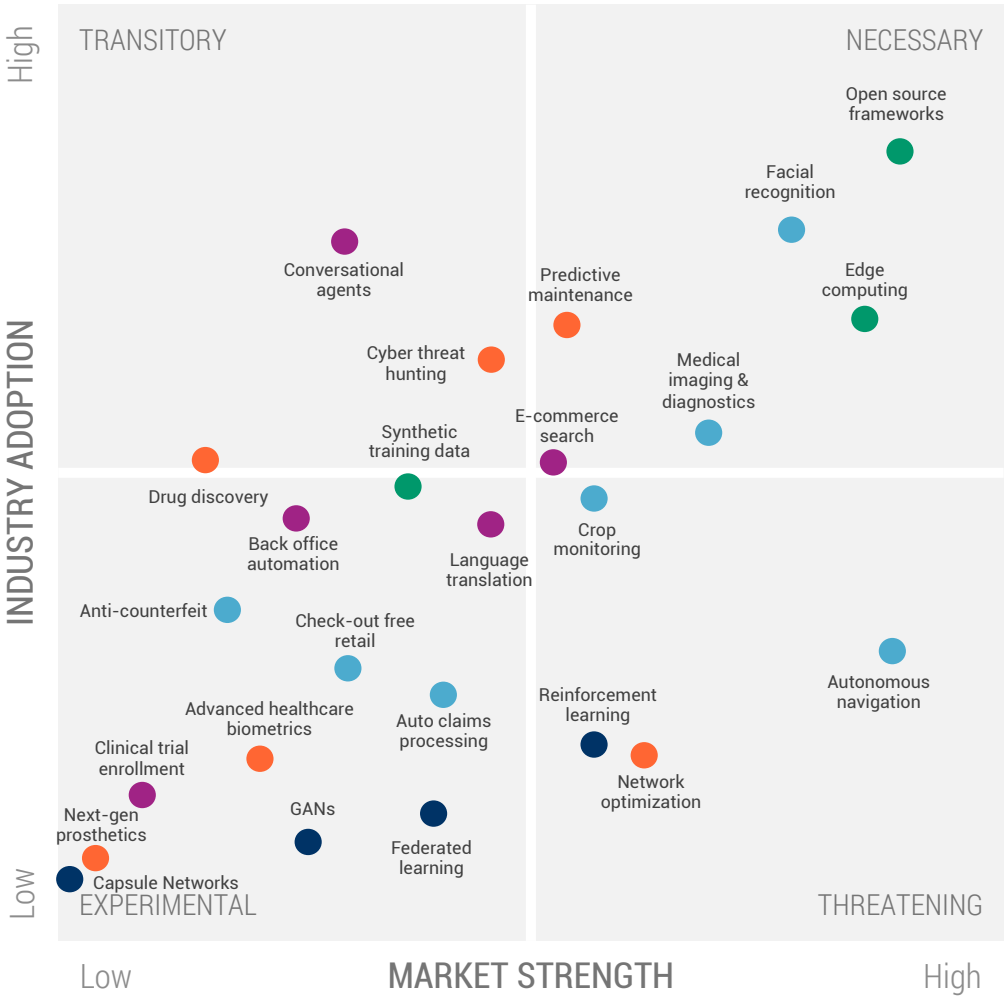

- Application: Computer vision
- Application: Natural language processing/synthesis
- Application: Predictive intelligence
- Architecture
- Infrastructure

NExTT Trends

We evaluate each of these trends using the CB Insights NExTT framework.

The NExTT framework educates businesses about emerging trends and guides their decisions in accordance with their comfort with risk.

NExTT uses data-driven signals to evaluate technology, product, and business model trends from conception to maturity to broad adoption.

The NExTT framework's 2 dimensions:

INDUSTRY ADOPTION (y-axis): Signals include momentum of startups in the space, media attention, customer adoption (partnerships, customer, licensing deals).

MARKET STRENGTH (x-axis): Signals include market sizing forecasts, quality and number of investors and capital, investments in R&D, earnings transcript commentary, competitive intensity, incumbent deal making (M&A, strategic investments).

NExTT framework's 2 dimensions

Industry Adoption (y axis)
Signals include:

momentum of startups in the space

media attention

customer adoption
(partnerships, customer, licensing deals)

Market Strength (x axis)
Signals include:

market sizing forecasts

earnings transcript commentary

quality and number of investors and capital

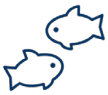

competitive intensity

investments in R&D

incumbent deal making
(M&A, strategic investments)

Necessary

OPEN-SOURCE FRAMEWORKS

The barrier to entry in AI is lower than ever before, thanks to open-source software.

Google open-sourced its TensorFlow machine learning library in 2015.

Open-source frameworks for AI are a two-way street: It makes AI accessible to everyone, and companies like Google, in turn, benefit from a community of contributors helping accelerate its AI research.

Source: GitHub

Hundreds of users contribute to TensorFlow every month on GitHub (a software development platform where users can collaborate).

Below are a few companies using TensorFlow, from Coca-Cola to eBay to Airbnb.

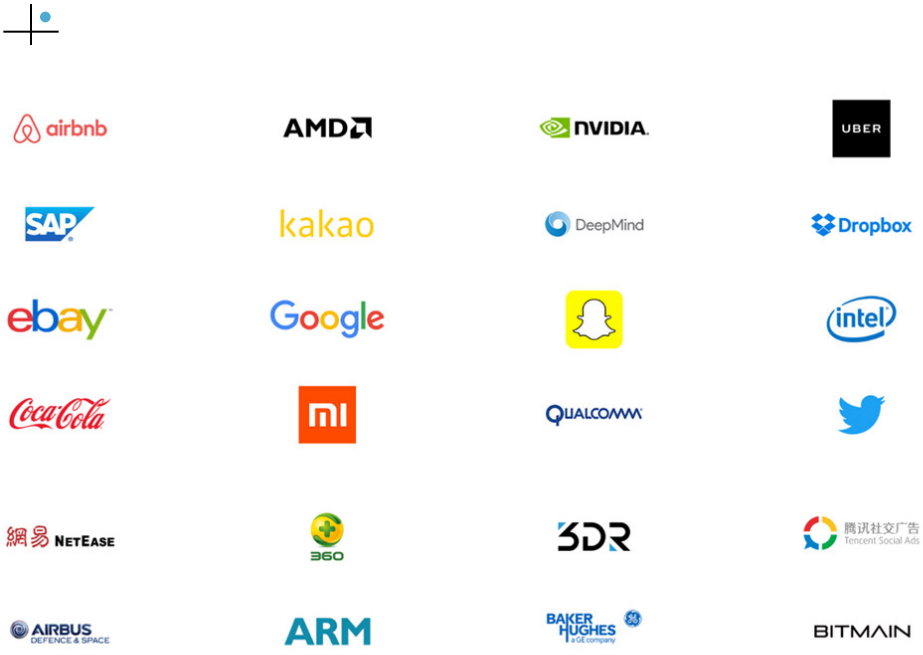

Source: TensorFlow

Facebook released Caffe2 in 2017, after working with researchers from Nvidia, Qualcomm, Intel, Microsoft, and others to create a “a lightweight and modular deep learning framework” that can extend beyond the cloud to mobile applications.

Facebook also operated PyTorch at the time, an open-source machine learning platform for Python. In May’18, Facebook merged the two under one umbrella to “combine the beneficial traits of Caffe2 and PyTorch into a single package and enable a smooth transition from fast prototyping to fast execution.”

The number of GitHub contributors to PyTorch have increased in recent months.

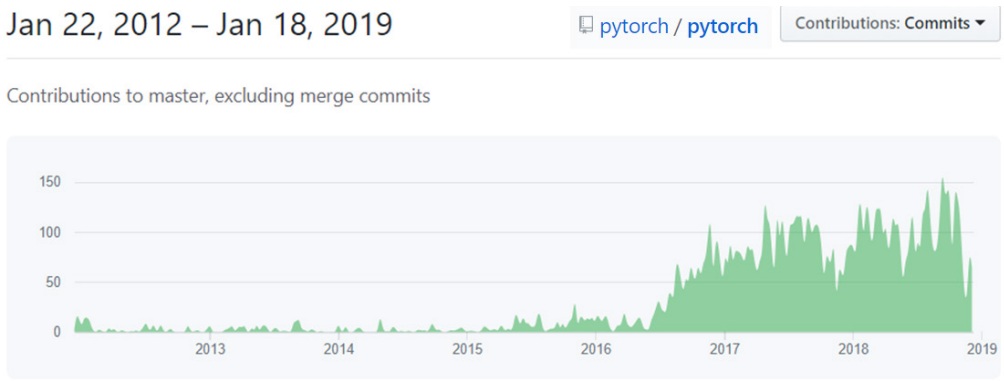

Source: GitHub

Theano is another open-source library from the Montreal Institute for Learning Algorithms (MILA). In Sep'17, leading AI researcher Yoshua Bengio announced an end to development on Theano from MILA as these tools have become so much more widespread.

"The software ecosystem supporting deep learning research has been evolving quickly, and has now reached a healthy state: open-source software is the norm; a variety of frameworks are available, satisfying needs spanning from exploring novel ideas to deploying them into production; and strong industrial players are backing different software stacks in a stimulating competition."

- YOSHUA BENGIO, IN A MILA ANNOUNCEMENT

A number of open-source tools are available today for developers to choose from, including Keras, Microsoft Cognitive Toolkit, and Apache MXNet.

EDGE AI

The need for real-time decision making is pushing AI closer to the edge.

Running AI algorithms on edge devices — like a smartphone or a car or even a wearable device — instead of communicating with a central cloud or server gives devices the ability to process information locally and respond more quickly to situations.

Nvidia, Qualcomm, and Apple, along with a number of emerging startups, are focused on building chips exclusively for AI workloads at the "edge."

From consumer electronics to telecommunications to medical imaging, edge AI has implications for every major industry.

For example, an autonomous vehicle has to respond in real-time to what's happening on the road, and function in areas with no internet connectivity. Decisions are time-sensitive and latency could prove fatal.

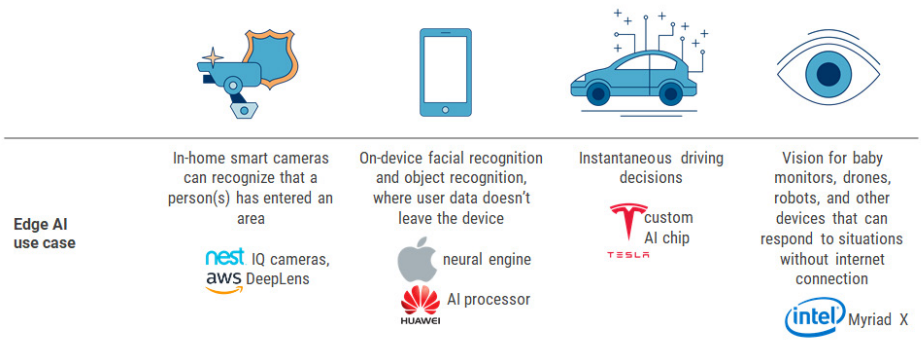

Big tech companies made huge leaps in edge AI between 2017-2018.

Apple released its A11 chip with a “neural engine” for iPhone 8, iPhone 8 Plus, and X in 2017, claiming it could perform machine learning tasks at up to 600 billion operations per second. It powers new iPhone features like Face ID, running facial recognition on the device itself to unlock the phone.

Qualcomm launched a $100M AI fund in Q4’18 to invest in startups “that share the vision of on-device AI becoming more powerful and widespread,” a move that it says goes hand-in-hand with its 5G vision.

As the dominant processor in many data centers, Intel has had to play catch-up with massive acquisitions. Intel released an on-device vision processing chip called Myriad X (initially developed by Movidius, which Intel acquired in 2016).

In Q4’18 Intel introduced the Intel NCS2 (Neural Compute Stick 2), which is powered by the Myriad X vision processing chip to run computer vision applications on edge devices, such as smart home devices and industrial robots.

The CB Insights earnings transcript analysis tool shows mentions of edge AI trending up for part of 2018.

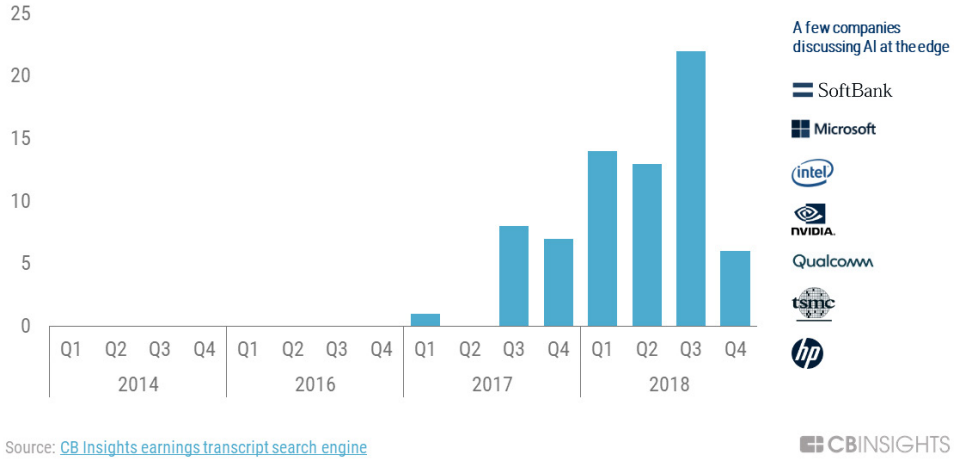

Microsoft said it introduced 100 new Azure capabilities in Q3'18 alone, "focused on both existing workloads like security and new workloads like IoT and edge AI."

Nvidia recently released the Jetson AGX Xavier computing chip for edge computing applications across robotics and industrial IoT.

While AI on the edge reduces latency, it also has limitations. Unlike the cloud, edge has storage and processing constraints. More hybrid models will emerge that allow intelligent edge devices to communicate with each other and a central server.

FACIAL RECOGNITION

From unlocking phones to boarding flights, face recognition is going mainstream.

When it comes to facial recognition, China's unapologetic push towards surveillance coupled with its AI ambitions have hogged the media limelight.

China's facial recognition trends up in news mentions

Number of news mentions of "China" and "facial recognition"

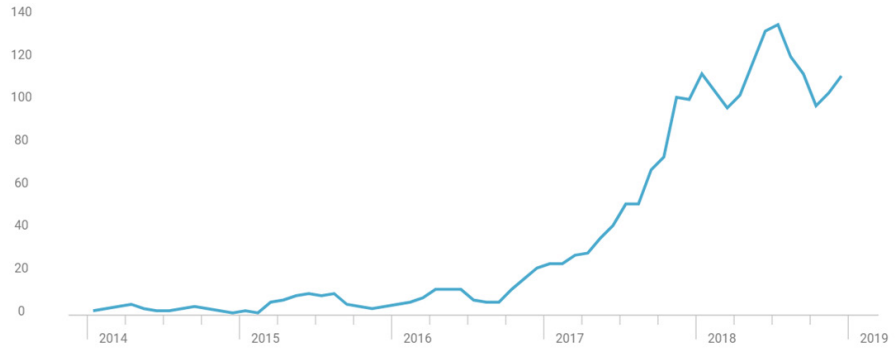

Source: CB Insights

CBINSIGHTS

As the government adds a layer of artificial intelligence to its surveillance, startups are playing a key role in providing the government with the underlying technology. A quick search on the CB Insights platform for face recognition startup deals in China reflect the demand for the technology.

Demand for facial recognition tech rises in China

Equity and non-equity deals, 2013 – 2018

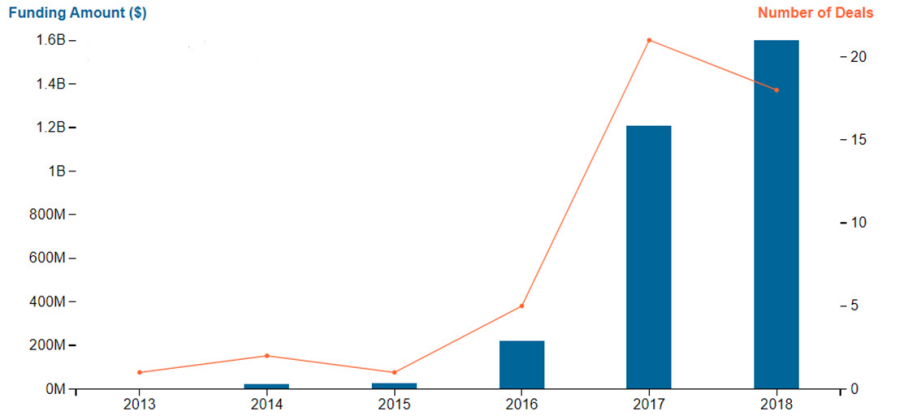

Source: CB Insights

CBINSIGHTS

Unicorns like SenseTime, Face++, and more recently, CloudWalk, have emerged from the country. (Here's our detailed report on China's surveillance efforts.)

But even in the United States, interest in the tech is surging, according to the CB Insights patent analysis tool.

US patent applications for facial recognition

Number of patents by date of filing, based on keyword searches in title and abstract

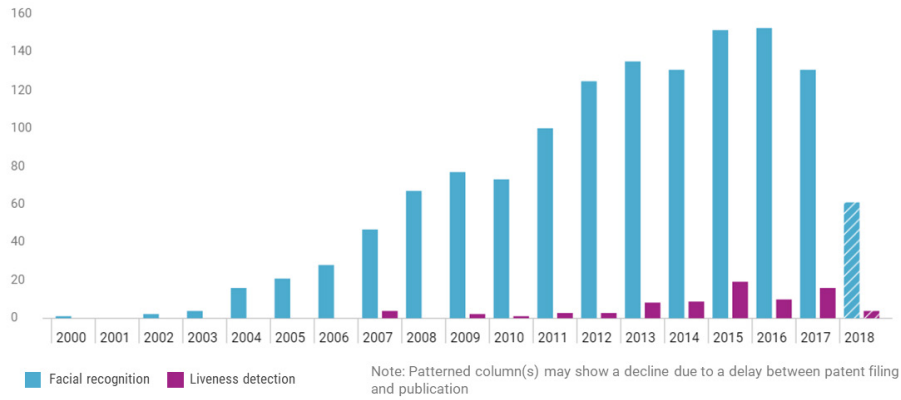

Note: Patterned column(s) may show a decline due to a delay between patent filing and publication

Source: CB Insights patent analytics

CBINSIGHTS

Apple popularized the tech for everyday consumers with the introduction of facial recognition-based login in iOS 10.

Amazon is selling its tech to law enforcement agencies.

Academic institutions like Carnegie Mellon University are also working on technology to help enhance video surveillance.

The university was granted a patent around “hallucinating facial features” — a method to help law enforcement agencies identify masked suspects by reconstructing a full face when only the periocular region of the face is captured. Facial recognition may then be used to compare the “hallucinated face” to images of actual faces to find ones with a strong correlation.

Source: USPTO

But the tech is not without glitches. Amazon was in the news for reportedly misidentifying some Congressmen as criminals.

Smart cameras outside a Seattle school were easily tricked by a WSJ reporter who used a picture of the headmaster to enter the premises, when the “smile to unlock feature” was temporarily disabled.

“Smile to unlock” and other such “liveness detection” methods offer an added layer of authentication.

For instance, Amazon was granted a patent that explores additional layers of security, including asking users to perform certain actions like “smile, blink, or tilt his or her head.”

These actions can then be combined with “infrared image information, thermal imaging data, or other such information” for more robust authentication.

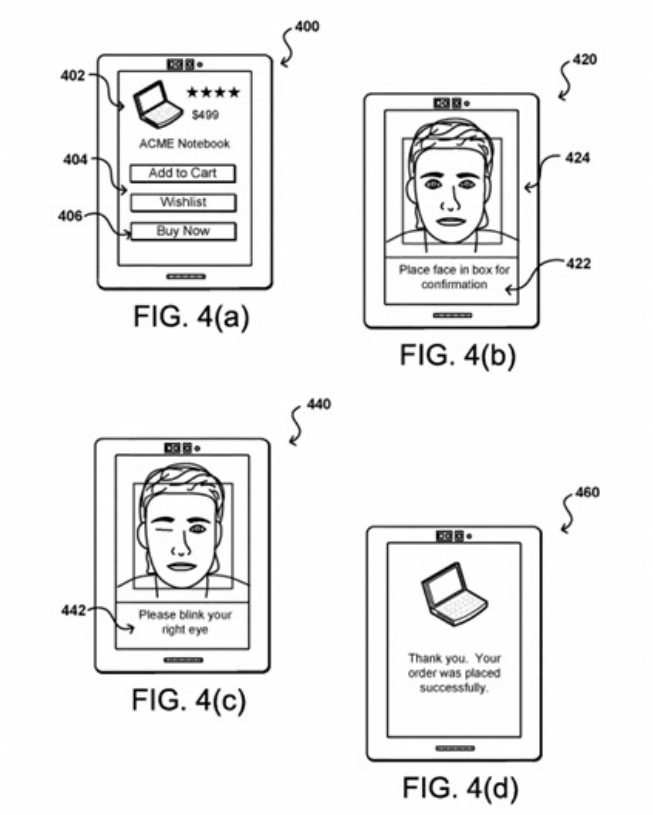

Source: USPTO

Early commercial applications are taking off in security, retail, and consumer electronics, and facial recognition is fast becoming a dominant form of biometric authentication.

MEDICAL IMAGING & DIAGNOSTICS

The FDA is greenlighting AI-as-a-medical-device.

In April 2018, the FDA approved AI software that screens patients for diabetic retinopathy without the need for a second opinion from an expert.

It was given a “breakthrough device designation” to expedite the process of bringing the product to market.

The software, IDx-DR, correctly identified patients with “more than mild diabetic retinopathy” 87.4% of the time, and identified those who did not have it 89.5% of the time.

IDx is one of the many AI software products approved by the FDA for clinical commercial applications in recent months.

Diabetic retinopathy

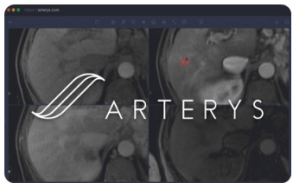

Liver and lung AI lesion

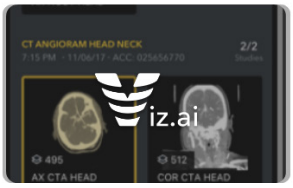

CT-scan analysis for strokes

The FDA cleared Viz LVO, a product from startup Viz.ai, to analyze CT scans and notify healthcare providers of potential strokes in patients. Post FDA clearance, Viz.ai closed a $21M Series A round from Google Ventures and Kleiner Perkins Caufield & Byers.

The FDA also cleared GE Ventures-backed startup Arterys for its Oncology AI suite initially focused on spotting lung and liver lesions.

Fast-track regulatory approval opens up new commercial pathways for over 80 AI imaging & diagnostics companies that have raised equity financing since 2014, accounting for a total of 149 deals.

Diagnostics is a major driver of health AI deals

Equity deals, 2014 – 2018

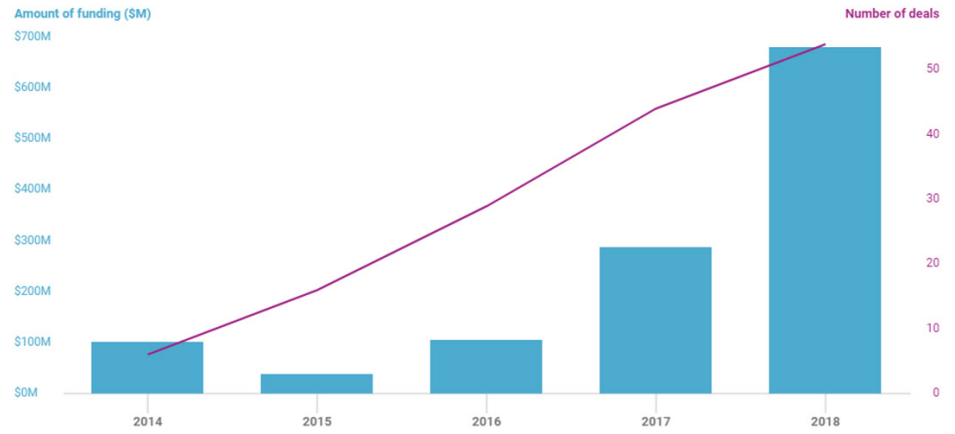

Source: CB Insights

CBINSIGHTS

On the consumer side, smartphone penetration and advances in image recognition are turning phones into powerful at-home diagnostic tools.

Startup Healthy.io's first product, Dip.io, uses the traditional urinalysis dipstick to monitor a range of urinary infections. Users take a picture of the stick with their smartphones, and computer vision algorithms calibrate the results to account for different lighting conditions and camera quality. The test detects infections and pregnancy-related complications.

Dip.io, which is already commercially available in Europe and Israel, was cleared by the FDA.

Apart from this, a number of ML-as-a-service platforms are integrating with FDA-approved home monitoring devices, alerting physicians when there is an abnormality.

PREDICTIVE MAINTENANCE

From manufacturers to equipment insurers, AI-IIoT can save incumbents millions of dollars in unexpected failures.

Field and factory equipment generate a wealth of data, yet unanticipated equipment failure is one of the leading causes of downtime in manufacturing.

A recent GE survey of 450 field service and IT decision makers found that 70% of companies are not aware of when equipment is due for an upgrade or maintenance, and that unplanned downtime can cost companies $250K/hour.

Predicting when equipment or individual components will fail benefits asset insurers, as well as manufacturers.

In predictive maintenance, sensors and smart cameras gather a continuous stream of data from machines, like temperature and pressure. The quantity and varied formats of real-time data generated make machine learning an inseparable component of IIoT. Over time, the algorithms can predict a failure before it occurs.

Dropping costs of industrial sensors, advances in machine learning algorithms, and a push towards edge computing have made predictive maintenance more widely available.

A leading indicator of interest in the space is the sheer number of big tech companies and startups here.

Deals to AI startups focused on industrials & energy

Equity deals, 2014 – 2018

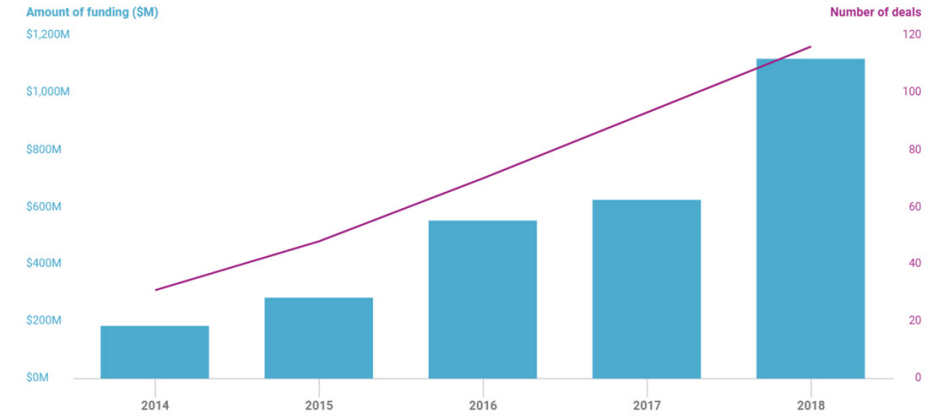

Source: CB Insights

CBINSIGHTS

Deals to AI companies focused on industrials and energy, which includes ML-as-a-service platforms for IIoT, are rising. Newer startups are competing with unicorns like C3 IoT and Uptake Technologies.

GE Ventures was an active investor here in 2016, backing companies including Foghorn Systems, Sight Machine, Maana, and Bit Stew Systems (which it later acquired). GE is a major player in IIoT, with its Predix analytics platform.

Competitors include Siemens and SAP, which have rolled out their own products (Mindsphere and Hana) for IIoT.

India's Tata Consultancy announced that it's launching predictive maintenance and AI-based solutions for energy utility companies. Tata claimed that an early version of its "digital twin" technology — replicating on-ground operations or physical assets in a digital format for monitoring them — helped a power plant save ~$1.5M per gigawatt per year.

Even big tech companies like Microsoft are extending their cloud and edge analytics solutions to include predictive maintenance.

E-COMMERCE SEARCH

Contextual understanding of search terms is moving out of the "experimental phase," but widespread adoption is still a long ways off.

Amazon has applied for over 35 US patents related to "search results" since 2002.

It has an exclusive subsidiary, A9, focused on product and visual search for Amazon. A9 has nearly 400 patent applications in the United States (not all of them related to search optimization).

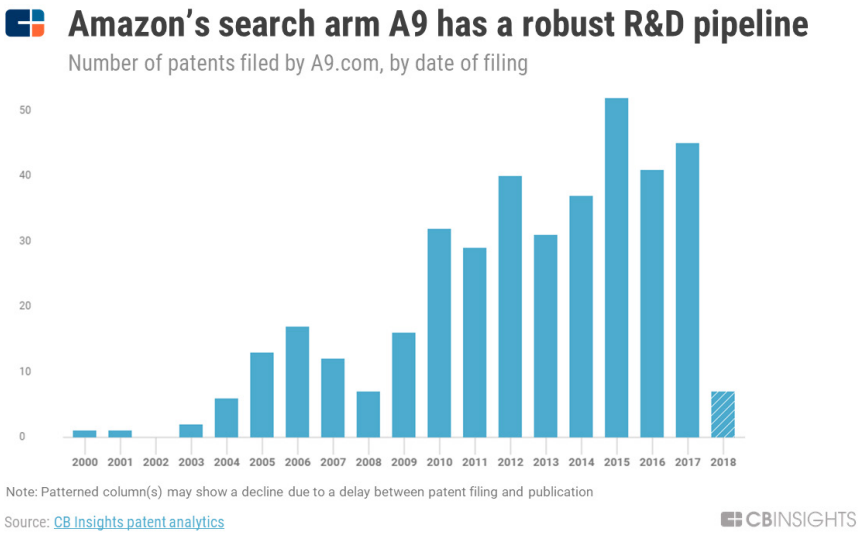

Some of the search-related patents include using convolutional neural networks to "determine a set of items whose images demonstrate visual similarity to the query image…" and using machine learning to analyze visual characteristics of an image and build a search query based on those.

Amazon is hiring for over 150 roles exclusively in its search division — for natural language understanding, chaos engineering, and machine learning, among other roles.

But Amazon's scale of operations and R&D in e-commerce search is the exception among retailers.

Few retailers have discussed AI-related strategies on earnings calls, and many haven't scaled or optimized their e-commerce operations.

But one of the earliest brands to do so was eBay.

The company first mentioned "machine learning" in its Q3'15 earnings calls. At the time, eBay had just begun to make it compulsory for sellers to write product descriptions, and was using machine learning to process that data to find similar products in the catalog.

Using proper metadata to describe products on a site is a starting point when using e-commerce search to surface relevant search results.

But describing and indexing alone is not enough. Many users search for products in natural language (like "a magenta shirt without buttons") or may not know how to describe what they're looking for.

This makes natural language for e-commerce search a challenge.

Early-stage SaaS startups are emerging, selling search technologies to third-party retailers.

Image search startup ViSenze works with clients like Uniqlo, Myntra, and Japanese e-commerce giant Rakuten. ViSenze allows in-store customers to take a picture of something they like at a store, then upload the picture to find the exact product online.

It has offices in California and Singapore, and raised a $10.5M Series B in 2016 from investors including the venture arm of Rakuten. It entered the Unilever Foundry in 2017, which allows startups in Southeast Asia to test pilot projects with its brands.

Another startup developing AI for online search recommendations is Israel-based Twiggle.

The Alibaba-backed company is developing a semantic API that sits on top of existing e-commerce search engines, responding to very specific searches by the buyer. Twiggle raised $15M in 2017 in a Series B round and entered the Plug and Play Accelerator last year.

Experimental

CAPSULE NETWORKS

Deep learning has fueled the majority of the AI applications today. It may now get a makeover thanks to capsule networks.

Google's Geoffrey Hinton, a pioneering researcher in deep learning, introduced a new concept called "capsules" in a paper way back in 2011, arguing that "current methods for recognizing objects in images perform poorly and use methods that are intellectually unsatisfying."

Those "current methods" Hinton referred to include one of the most popular neural network architectures in deep learning today, known as convolutional neural networks (CNN). CNN has particularly taken off in image recognition applications. But CNNs, despite their success, have shortcomings (more on that below).

Hinton published 2 papers during 2017-2018 on an alternative concept called "capsule networks," also known as CapsNet — a new architecture that promises to outperform CNNs on multiple fronts.

Without getting into the weeds, CNNs fail when it comes to precise spatial relationships. Consider the face below. Although the relative position of the mouth is off with respect to other facial features, a CNN would still identify this as a human face.

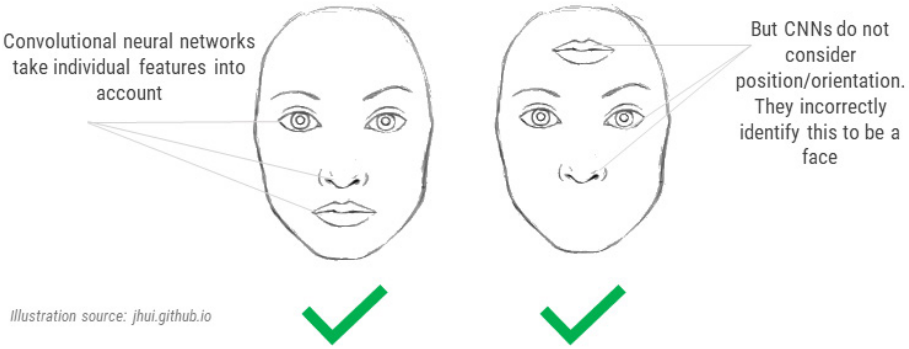

Although there are methods to mitigate the above problem, another major issue with CNNs is the failure to understand new viewpoints.

“Now that convolutional neural networks have become the dominant approach to object recognition, it makes sense to ask whether there are any exponential inefficiencies that may lead to their demise. A good candidate is the difficulty that convolutional nets have in generalizing to novel viewpoints.”

— PAPER ON DYNAMIC ROUTING BETWEEN CAPSULES

For instance, a CapsNet does a much better job of identifying the images of toys in the first and second rows as belonging to the same object, only taken from a different angle or viewpoint. CNNs would require a much larger training dataset to identify each orientation.

(The images above are from a database called smallNORB which contains grey-scale images of 50 toys belonging to 1 of 5 categories: four-legged animals, human figures, airplanes, trucks, and cars. Hinton's paper found that CapsNets reduced the error rate by 45% when tested on this dataset compared to other algorithmic approaches.)

Hinton claims that capsule networks were tested against some sophisticated adversarial attacks (tampering with images to confuse the algorithms) and were found to outperform convolutional neural networks.

Hackers can introduce small variations to fool a CNN. Researchers at Google and OpenAI have demonstrated this with several examples.

One of the more popular examples CapsNet was tested against is from a 2015 paper by Google's Ian Goodfellow and others. As can be seen below, a small change that is not readily noticeable to the human eye means the image results in a neural network identifying a panda as a gibbon, a type of ape, with high confidence.

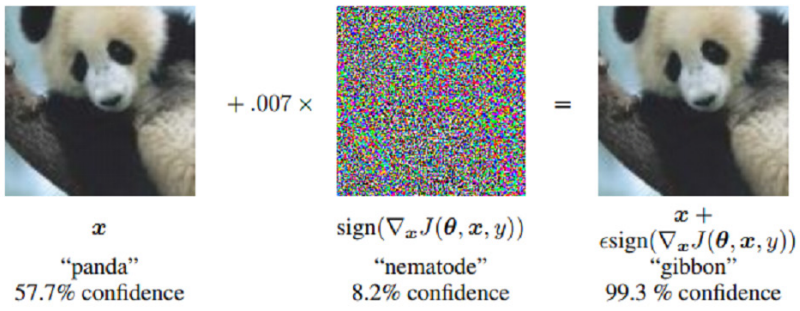

Research into capsule networks is in its infancy, but could challenge current state-of-the-art approaches to image recognition.

NEXT-GEN PROSTHETICS

Very early-stage research is emerging, combining biology, physics, and machine learning to tackle one of the hardest problems in prosthetics: dexterity.

DARPA has spent millions of dollars on its advanced prosthetics program, which it started in 2006 with John Hopkins University to help wounded veterans. But the problem is a complex one to tackle.

For instance, giving amputees the ability to move individual fingers in a prosthetic arm, decoding brain and muscle signals behind voluntary movements, and translating that into robotic control all require a multi-disciplinary approach.

As Megan Molteni explained in an article for Wired last year, take a simple example of playing the piano. After repeated practice, playing a chord becomes "muscle memory," but that's not how prosthetic limbs work.

More recently, researchers have started using machine learning to decode signals from sensors on the body and translate them into commands that move the prosthetic device.

John Hopkins' Applied Physics Labs has an ongoing project on neural interfaces for prosthetics using "neural decoding algorithms" to do just that.

In June last year, researchers from Germany and Imperial College London used machine learning to decode signals from the stump of the amputee and power a computer to control the robotic arm. The research on the "brain-machine interface" was published in Science Robotics.

Other papers explore intermediary solutions like using myoelectric signals (electric activity of muscles near the stump) to activate a camera, and running computer vision algorithms to estimate the grasp type and size of the object before them.

Further highlighting the AI community's interest in the space, the "AI for Prosthetics Challenge" was one of the competition tracks in NeurIPS'18 (a leading, annual machine learning conference).

The 2018 challenge was to predict the performance of a prosthetic leg using reinforcement learning (more on reinforcement learning in the following sections of this report). Researchers use an open-source software called OpenSim which simulates human movement.

The previous year's focus was "Learning to Run," which saw 442 participants attempting to teach AI how to run, with sponsors including AWS, Nvidia, and Toyota.

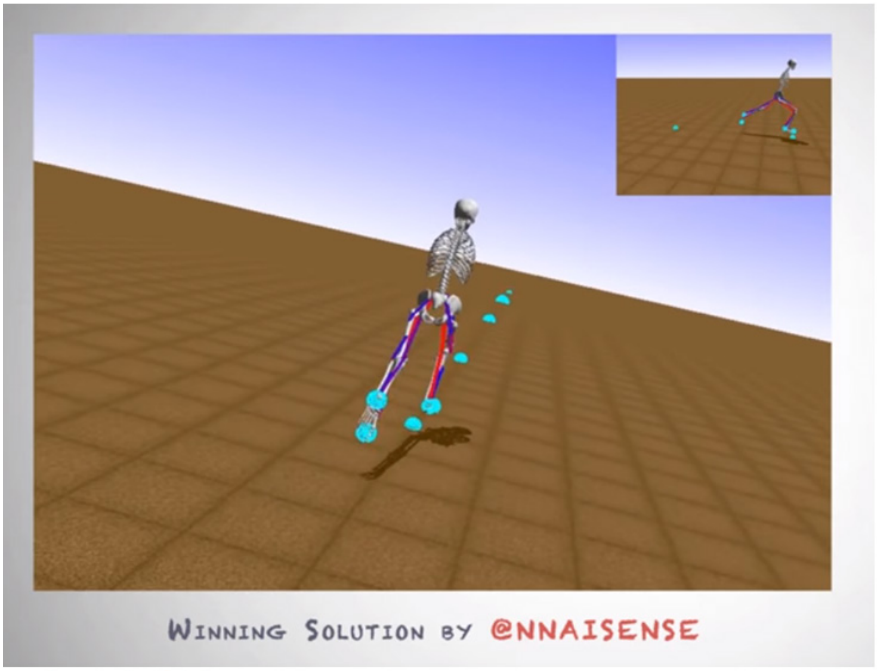

Source: http://osim-rl.stanford.edu/

CLINICAL TRIAL ENROLLMENT

One of the biggest bottlenecks in clinical trials is enrolling the right pool of patients. Apple might be able to solve this issue.

Interoperability — the ability to share information easily across institutions and software systems — is a one of the biggest issues in healthcare, despite efforts to digitize health records.

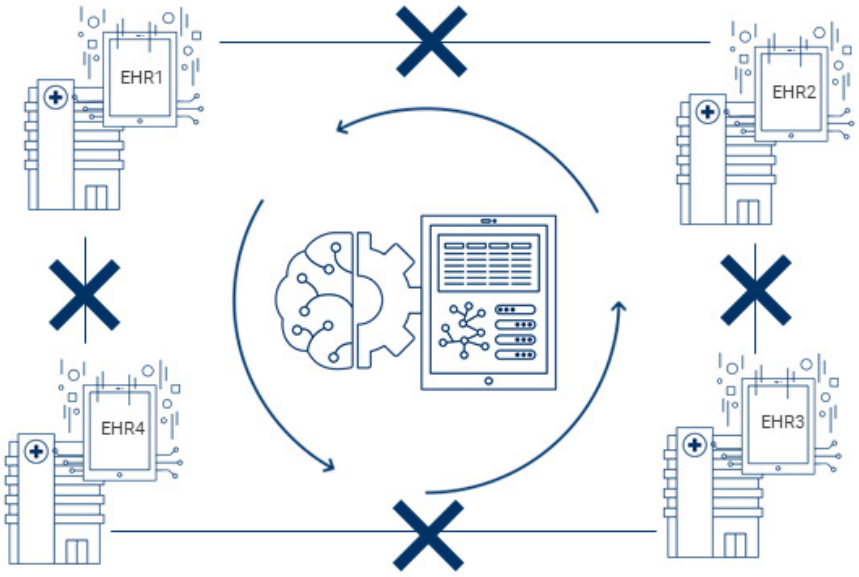

This is particularly problematic in clinical trials, where matching the right trial with the right patient is a time-consuming and challenging process for both the clinical study team and the patient.

For context, there are over 18,000 clinical studies that are currently recruiting patients in the US alone.

Patients may occasionally get trial recommendations from their doctors if a physician is aware of an ongoing trial.

Otherwise, the onus of scouring through ClinicalTrials.Gov — a comprehensive federal database of past and ongoing clinical trials — falls on the patient.

An ideal AI solution would be artificial intelligence software that extracts relevant information from a patient's medical records, compares it with ongoing trials, and suggests matching studies.

Few startups are working with clients directly in the clinical trials space. The biggest barriers to entry for smaller startups streamlining clinical trials are that the technologies are relatively new and the industry is slow to adapt.

Tech giants like Apple, however, have seen success in bringing on partners for their healthcare-focused initiatives.

Apple is changing how data flows in healthcare and is opening up new possibilities for AI, specifically around how clinical study researchers recruit and monitor patients.

Since 2015, Apple has launched two open-source frameworks — ResearchKit and CareKit — to help clinical trials recruit patients and monitor their health remotely.

The frameworks allow researchers and developers to create medical apps to monitor people's daily lives, removing geographic barriers to enrollment.

For example, nearly 10,000 people use the mPower app, which provides exercises like finger tapping and gait analysis to study patients with Parkinson's disease who have consented to share their data with the broader research community.

Researchers at Duke University developed an Autism & Beyond app that uses the iPhone's front camera and facial recognition algorithms to screen children for autism.

Apple is also working with popular EHR vendors like Cerner and Epic to solve interoperability problems.

In January 2018, Apple announced that iPhone users would have access to all their electronic health records from participating institutions on their iPhone's Health app.

Called "Health Records," the feature is an extension of what AI healthcare startup Gliimpse was working on before it was acquired by Apple in 2016.

In an easy-to-use interface, users can find all the information they need on allergies, conditions, immunizations, lab results, medications, procedures, and vitals.

In June 2018, Apple rolled out a Health Records API for developers.

Users can now choose to share their data with third-party applications and medical researchers, opening up new opportunities for disease management and lifestyle monitoring.

The possibilities are seemingly endless when it comes to using AI and machine learning for early diagnosis, enrolling the right pool of patients, and even driving decisions in drug design.

GENERATIVE ADVERSARIAL NETWORKS

Two neural networks trying to outsmart each other are getting very good at creating realistic images.

Can you identify which of these images are fake?

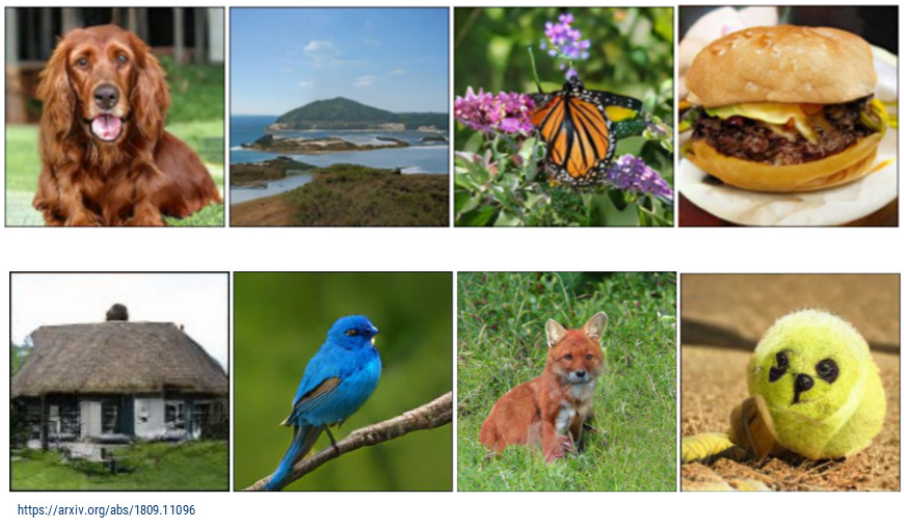

https://arxiv.org/abs/1809.11096

The answer is all of the above. Each of these highly realistic images were created by generative adversarial networks, or GANs.

(Note: the bottom right image represents a “class leakage” — where the algorithm possibly confused properties of a dog with a ball — and created a “dogball”)

GAN, a concept introduced by Google researcher Ian Goodfellow in 2014, taps into the idea of “AI versus AI.” There are two neural networks: the generator, which comes up with a fake image (say a dog for instance), and a discriminator, which compares the result to real-world images and gives feedback to the generator on how close it is to replicating a realistic image.

This forms a constant feedback loop between two neural networks trying to outsmart each other.

The images above are from a Sept'18 paper by Andrew Brock, an intern at Google DeepMind, published along with other DeepMind researchers. They trained GANs on a very large scale dataset to create "BigGANs."

One of the challenges Brock and team encountered with BigGANs: A spider, for example, has "lots of legs." But how many is "lots"?

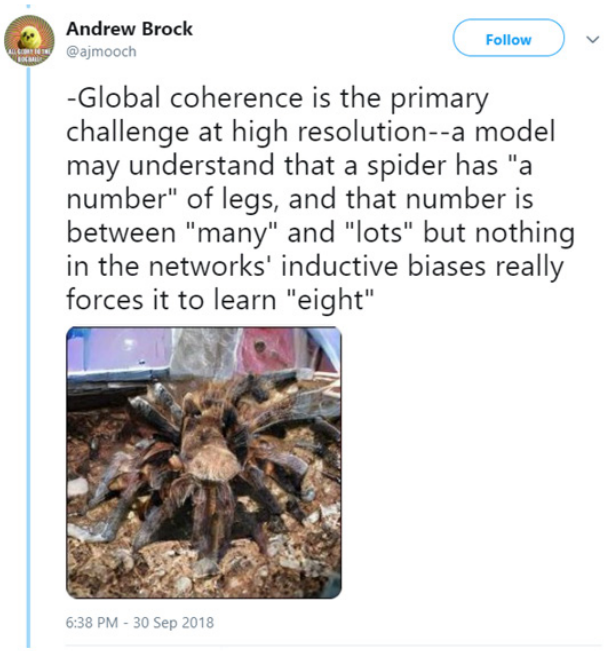

The primary challenge to scaling large-scale projects like GANs, however, is computational power. Here's an excerpt from FastCompany, with a rough estimation of the amount of computing power that went into this research:

> These experiments have environmental implications as well. Brock used 512 of Google's Tensor Processing Units (or TPU) to generate his 512 pixel images, and he says his experiments generally run for between 24 and 48 hours. If each TPU uses about 200 watts in an hour of computation, then a single one of Brock's 512 pixel experiments could be using between 2,450 and 4,915 kilowatt hours. That's the equivalent of the electricity that the average American household uses in just under six months.

For GANs to scale, hardware for AI has to scale in parallel.

Brock's is not the only GAN-related paper published in recent months.

Using GANs, researchers from Lancaster University in the UK, Northwest University in the China, and Peking University in China developed a captcha solver.

The paper demonstrated that GANs can crack text-based captchas in just 0.05 seconds using a desktop GPU, with a relatively higher success rate compared to previous methods.

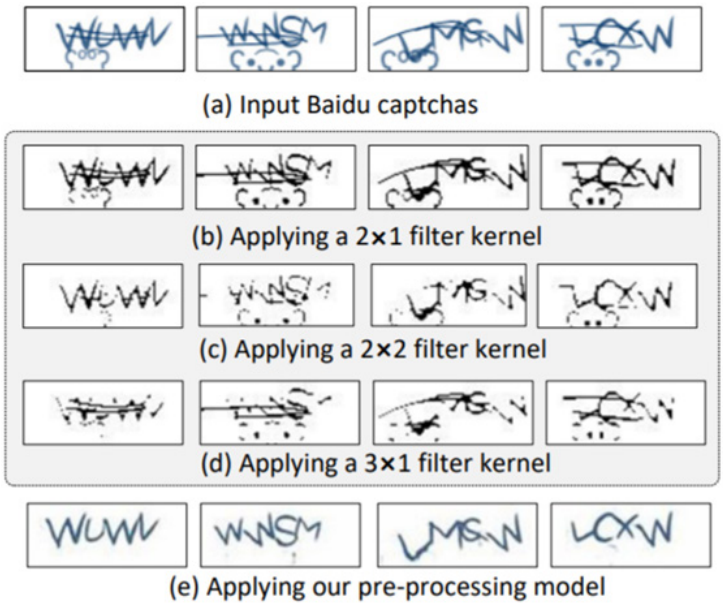

Source: lancaster.ac.uk/staff/wangz3/publications/ccs18.pdf

Researchers at CMU used GANs for “face-to-face” translation in this iteration of “deepfake” videos. In the deepfake example below, John Oliver turns into Stephen Colbert:

Source: http://www.cs.cmu.edu/~aayushb/Recycle-GAN/

Researchers at the Warsaw University of Technology developed a ComixGAN framework to turn videos into comics using GANs.

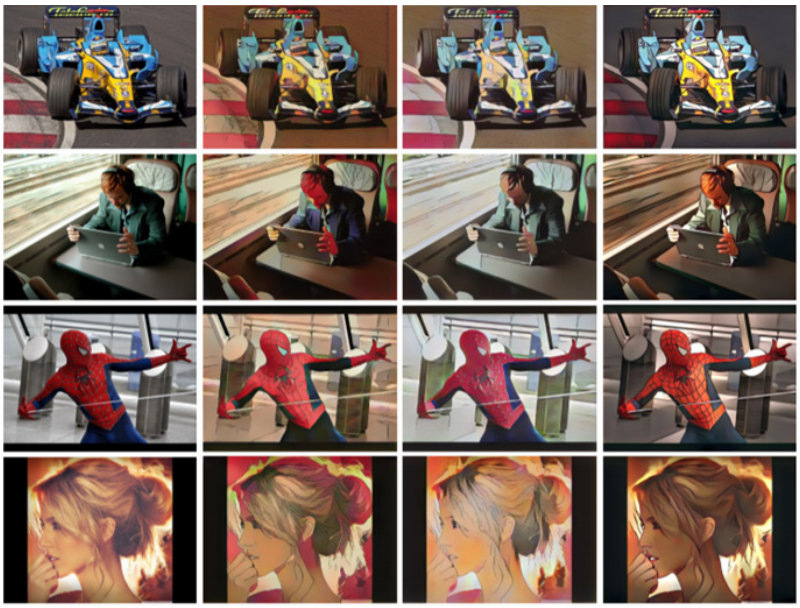

Figure 8: Results for different comixification approaches. Columns: a) content images, b) original CartoonGAN model (Hayao), c) original CartoonGAN model (Hosoda), d) ComixGAN.

Source: https://arxiv.org/pdf/1812.03473.pdf

Art auction house Christie's sold its first ever GAN-generated painting for a whopping $432,500.

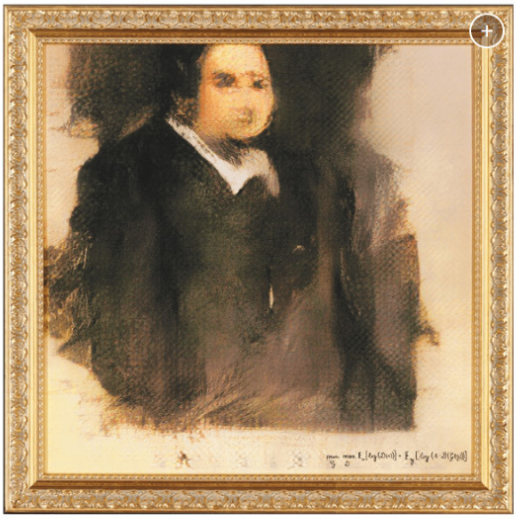

Portrait of Edmond Belamy, 2018, created by GAN (Generative Adversarial Network). Sold for $432,500 on 25 October at Christie's in New York. Image © Obvious

And in a more recent paper on GANs, Nvidia researchers used a "style-based generator" to create hyper-realistic images.

Source: https://arxiv.org/pdf/1812.04948.pdf

GANs aren't just for fun experiments. The approach also has serious implications, including fake political videos and morphed pornography. The Wall Street Journal is already training its researchers to spot deepfake videos.

As the research scales, it will change the future of news, media, art, and even cybersecurity. GANs are already changing how we train AI algorithms (more on this in the following section on "synthetic training data.")

FEDERATED LEARNING

The new approach aims to protect privacy while training AI with sensitive user data.

Our daily interaction with smartphones and tablets — from the choice of words we use in messaging to the way we react to photos — generates a wealth of data.

Training AI algorithms using our unique local datasets can vastly improve their performance, such as more accurately predicting the next word you're going to type into your keyboard.

As researchers from Google explain in a 2017 paper, "the use of language in chat and text messages is generally much different than standard language corpora, e.g., Wikipedia and other web documents; the photos people take on their phone are likely quite different than typical Flickr photos."

But this user data is also personal and privacy sensitive.

Google's federated learning approach aims to use this rich dataset, but at the same time protect sensitive data.

In a nutshell, your data stays on your phone. It is not sent to or stored in a central cloud server. A cloud server sends the most updated version of an algorithm — called the "global state" of the algorithm — to a random selection of user devices.

Your phone makes improvements and updates to the model based on your localized data. Only this update (and updates from other users) are sent back to the cloud to improve the "global state" and the process repeats itself.

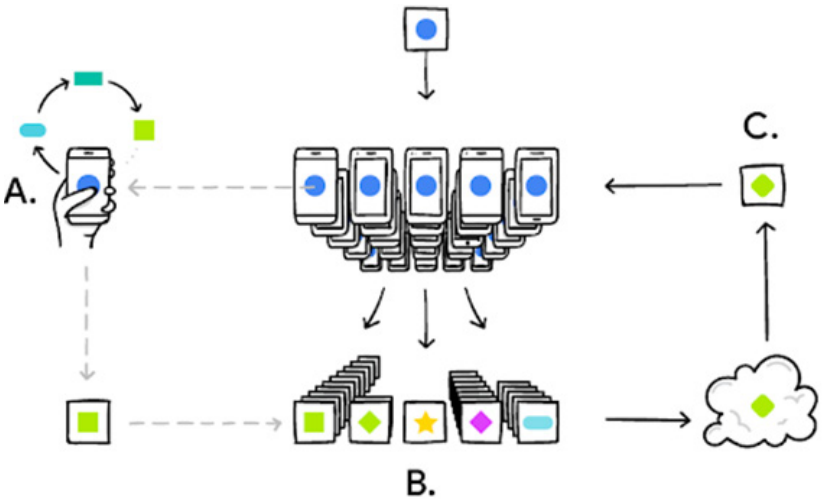

Source: Google AI blog

Google is testing federated learning in its Android keyboard called Gboard.

Note that the mechanism of aggregating individual updates from each node is not the novelty here. There are algorithms that do that already.

But unlike other distributed algorithms, the federated learning approach takes into account two important characteristics of the dataset:

- **Non-IID:** Data generated on each phone (or other device) is unique based on each person's usage of the device. And so these datasets are not "Independent and identically distributed (IID)" — a common assumption made by other distributed algorithms for the sake of statistical inference, but not reflective of practical real-world scenarios.
- **Unbalanced:** Some users are more actively engaged with an app than others, naturally generating more data. As a result, each phone, for instance, will have varying amounts of training data.

Firefox tested out federated learning to rank suggestions that appear when a user starts typing into the URL bar, calling it “one of the very first implementations [of federated learning] in a major software project.”

In another application of federated learning, Google Ventures-backed AI startup OWKIN, which is focused on drug discovery, is using the approach to protect sensitive patient data. The model allows different cancer treatment centers to collaborate without patients’ data ever leaving the premises, according to investor Otium Venture.

ADVANCED HEALTHCARE BIOMETRICS

Using neural networks, researchers are starting to study and measure atypical risk factors that were previously difficult to quantify.

Analysis of retinal images and voice patterns using neural networks could potentially help identify risk of heart disease.

Researchers at Google used a neural network trained on retinal images to find cardiovascular risk factors, according to a paper published in Nature this year.

The research found that not only was it possible to identify risk factors such as age, gender, and smoking patterns through retinal images, it was also "quantifiable to a degree of precision not reported before."

Similarly, the Mayo Clinic partnered with Beyond Verbal, an Israeli startup that analyzes acoustic features in voice, to find distinct voice features in patients with coronary artery disease (CAD). The study found 2 voice features that were strongly associated with CAD when subjects were describing an emotional experience.

Recent research from startup Cardiogram suggests "heart rate variability changes driven by diabetes can be detected via consumer, off-the-self wearable heart rate sensors" using deep learning. One algorithmic approach showed 85% accuracy in detecting diabetes from heart rate.

A more futuristic use case is passive monitoring of healthcare biometrics.

In January 2018, a Google patent was published with an ambitious vision

for analyzing cardiovascular function from a person's skin color or skin displacement.

The sensors might even be positioned (per the patent's illustrations) in a "sensing milieu" in a patient's bathroom.

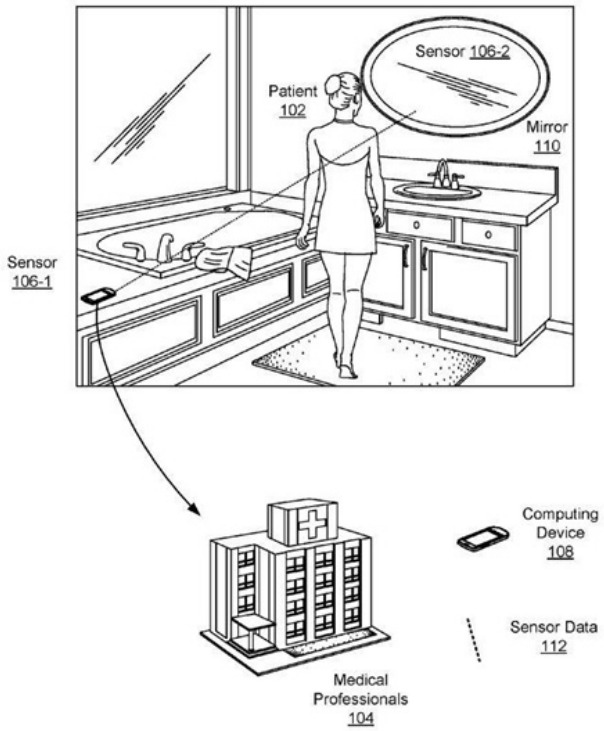

Source: USPTO

By recognizing skin color changes at the wrist and cheek, for example, and "comparing the times [of measurement] and distance between these regions," the system could calculate a "pulse-wave velocity (PWV)." The velocity information could then be used to determine cardio-health metrics such as arterial stiffness or blood pressure.

"Machine learning could be applied to create a patient specific model for estimating blood pressure from PWV," according to the patent.

Amazon applied for a similar patent for passive monitoring in 2014,

which was later granted in 2017. It combines recognition of facial features (using neural nets or other algorithmic approaches) with heart rate analysis.

For example, algorithms can track color changes in two areas of the face, like regions near the eyes and cheek, using that data to calculate heart rate detection.

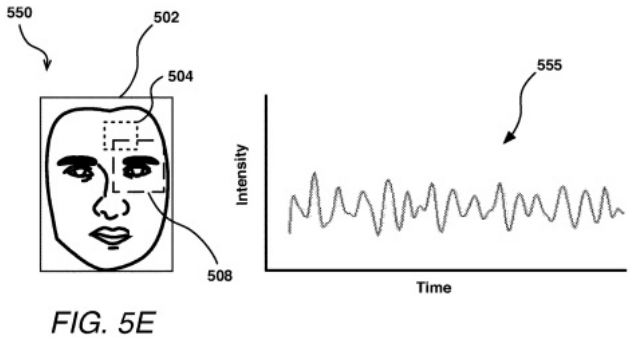

FIG. 5E

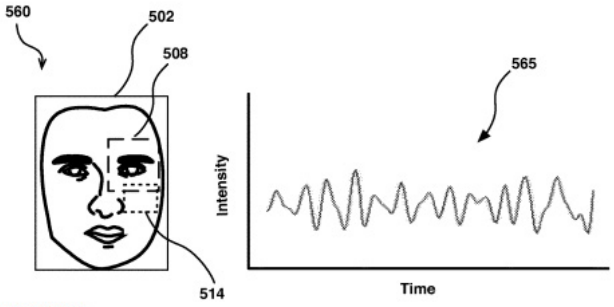

Source: USPTO

AI's ability to find patterns will continue to pave the way for new diagnostic methods and identification of previously unknown risk factors.

AUTO CLAIMS PROCESSING

Insurers and startups are beginning to use AI to compute a car owner's "risk score," analyze images of accident scenes, and monitor driver behavior.

China's Ant Financial, an Alibaba affiliate, uses deep-learning algorithms for image processing in its "accident processing system."

Currently, car owners or drivers take their vehicles to an "adjuster," a person who inspects the damage to the vehicle and logs the details, which are then sent to the auto insurance company.

Advances in image processing are now allowing people to take a picture of the vehicle and upload it to Ant Financial. Neural networks then analyze the image and automate the damage assessment.

Another approach Ant is taking is to create a risk profile of the driver to influence the actual pricing model of auto insurance.

"The development of technologies such as Big Data and artificial intelligence enables insurance companies to further leverage the consumer data and analyze the probable risk exposure of vehicle owners. Therefore, risk factors for auto insurance can shift from a "car-oriented" approach to a "car/owner combination."

— ALIBABA CLOUD BLOG

Alibaba introduced something called “Auto Insurance Points,” using machine learning to calculate a car owner’s risk score based on factors such as credit history, spending habits, and driving habits, among other things.

Smaller startups are also getting into insurance and claims processing but adopting a different approach.

Nexar, for instance, incentivizes drivers to use their smartphones as a dashcam and upload the footage to the Nexar app. In return, owners get a discount on their insurance premiums.

The app uses computer vision algorithms to monitor road conditions, driver behavior, and accidents. It also offers a “crash recreation” feature to reconstruct and analyze the circumstances in which accidents take place, and works with insurance clients to process claims.

UK-based Tractable allows insurers to upload an image of the damage and an estimate into its claims management platform. The “AI Review” feature compares this with thousands of images to adjust the price accordingly.

Interestingly, Tractable is targeting other players in the ecosystem as well, such as car repairers, appraisers, vendors, and car hire companies.

ANTI-COUNTERFEITING

Fakes are getting harder to spot, and online shopping makes it easier than ever to buy fake goods. To fight back, brands and pawnbrokers are beginning to experiment with AI.

From drugs to handbags to smartphones, counterfeiting is a problem that affects all types of retail.

Some product imitations look so authentic that they are classified as "super fakes."

China's rapidly growing e-commerce platform Pinduoduo mentioned "counterfeit" 11 times in its Q3'18 earnings call, describing "a very hard fight against counterfeit goods and … problematic merchants."

"In 2017, we…proactively removed a total of 10.7 million problematic products and blocked 40 million links that…raised infringement issues…We have also partnered with over 400 brands to work together on combating counterfeit."

— COLIN HUANG, FOUNDER AND CEO OF PINDUODUO

Brands are fighting the war against fakes on two fronts:

- In the online world, identifying and removing online listings that infringe on brand trademarks like brand name, logo, and slogans
- In the physical world, identifying fake goods like luxury handbags that are rip offs

Online counterfeiting is vast and complex in scope and scale.

E-commerce giant Alibaba, which has been under some fire for not doing enough to counter fake goods on its sites, reported that it's using deep learning to continuously scan its platform for IP infringements. It uses image recognition to identify characters in images, coupled with semantic recognition, possibly to monitor brand names or slogans in images of products listed on its sites.

Counterfeiters use keywords and images very similar to the original brand listing to sell fake goods on fake websites, fake goods on legitimate marketplaces, and promote fake goods on social media sites like Instagram.

When one listing is taken down, counterfeiters may repost the same fake product with a different string of keywords.

Barcelona-based startup Red Points is using machine learning to scan websites for potential infringements and find patterns in the choice of keywords counterfeiters use. It boasts clients in the cosmetics, luxury watch, home goods, and apparel industries, including MVMT, DOPE, and Paul Hewitt.

Spotting fakes is trickier and more manual in the physical world.

When a seller posts a second-hand luxury handbag for sale, or goes to a pawnbroker to trade it, the verification process usually involves an authentication expert physically examining the bag, including the make, material, and stitching pattern.

Here's how much eBay and others charge to authenticate one luxury handbag using identification experts.

Markup for authenticating a $1,500 handbag

20%

ebay POSHMARK

30-45%

The RealReal

Source: eBay Authenticate pricing comparison

But with the rise of "super fakes" or "triple-A fakes," it's becoming nearly impossible to tell the difference with the naked eye.

Building a database of fake and authentic goods, extracting their features, and training an AI algorithm to tell the difference is a cumbersome process.

Startup Entrupy worked with authentication experts to build a database of fake vs. real goods for training its algorithms for 2 years. The process is harder for rare vintage luxury goods.

Entrupy developed a portable microscope that attaches to a smartphone. When users take and upload a picture of the product (handbag, watch, etc), AI algorithms analyze microscopic signatures that are unique to each product, and verify it against a database of known and authentic products.

The database is growing, but there isn't a complete set products out in the market. A paper published by Entrupty highlights some other operating assumptions and limitations.

The key idea is that objects manufactured using standard or prescribed methods will have visually similar characteristics, compared to the manufacturing process a counterfeiter would use (non-standardized, inexpensive mass production). Secondly, the tech may not work for things like electronic chips that are nano-fabricated (variations at a scale that Entrupy's microscope cannot detect).

Cypheme is taking a different approach. Its ink-based technology can be used as a sticker on the product, or directly printed onto labels and packaging.

Random ink prints verifiable by neural nets

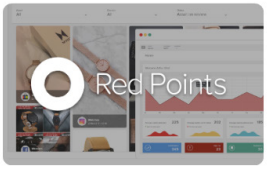

Online verification with machine learning

Handheld microscopic device

Nikkei Asian Review detailed the tech in an interview with the CEO: A random pattern is generated from a drop of ink, the pattern is surrounded by another circle of orange ink that Cypheme claims is proprietary to the company and impossible to replicate, then each unique pattern is associated with a specific product on a database.

It uses a smartphone camera and neural networks for pattern recognition to verify the ink pattern for the specific product against its database.

This means Cypheme has to work directly with brand manufacturers to make sure products are shipped with the tracing ink. It recently entered into a partnership with AR Packaging, a leading packaging company in Europe working with food brands like Unilever and Nestle.

While printing ink on packaging is efficient for tracking an item from the manufacturing plant and along the distribution chain, the tech doesn't work for secondhand purchase authentication. For instance, a buyer may remove Cypheme's sticker from the packaging of a luxury watch, and decide to resell it at a broker shop or online. In this case, verifying authenticity is not possible unless the printing is part of the product itself.

The solution for luxury brands and other high-stake retailers, moving forward, may be to identify or add unique fingerprints to physical goods at the site of manufacturing and track it through the supply chain.

CHECKOUT-FREE RETAIL

Entering a store, picking what you want, and walking out almost "feels" like shoplifting. AI could make actual theft a thing of the past and check-out free retail much more common.

Amazon Go did away with the entire checkout process, allowing shoppers to grab items and walk out.

Amazon has no public plans to sell its tech-as-a-service to other retailers yet, and has been tight-lipped about the operations, success, and pain points — only revealing that it uses sensors, cameras, computer vision, and deep learning algorithms. It has denied using facial recognition algorithms.

Startups like Standard Cognition and AiFi have seized the opportunity, stepping in to democratize Amazon Go for other retailers.

A challenge for grab-and-go stores is charging the right amount to the right shopper.

Loss of inventory due to shoplifting and paperwork error, among other things, cost US retailers around $47B in 2017, according to the National Retail Federation.

"Stealing is buying," Steve Gu, co-founder and CEO of startup AiFi, said in an interview with The AI Podcast, discussing the technology behind grab-and-go stores.

So far, Amazon Go is the only successful commercial deployment, but the the parameters of success are tightly controlled.

The chance of someone shoplifting is minimized when you control who enters the store, and automatically charge them.

Amazon already has an established base of Prime members. All the Go stores so far have been restricted to members, with other retail operations like the Kindle store, which is open to the general public, still relying on a manual checkout process.

Smaller bodegas, convenience stores, and even several established supermarkets have to build that membership base from scratch.

Steve Gu hinted in the same podcast that there could be a "grab-and-go" section for people willing to download the app, and a separate checkout line for those who don't want to.

It's not clear how a store's infrastructure would support both.

That still leaves the issue of point-of-sale inventory shrinkage such as incorrectly billed items or POS theft. China's Yitu Technology and Toshiba, with its intelligent camera for checkout, are some of the companies separately working on the shrinkage problem.

The complexity of preventing theft depends on the size and scale of operations, and type of products on the shelves.

Amazon Go stores are only about 1,800 to 3,000 sq. ft, and use hundreds of cameras covering nearly every inch of ceiling space. In comparison, traditional supermarkets can be 40,000 sq. ft. or more.

Go, which uses weight sensors on shelves in addition to cameras for visual recognition, currently only offers a limited selection of items, like prepared and packaged meal kits.

Some things to consider are how floor space will be utilized, especially in densely packed supermarkets, to ensure cameras are optimally placed to track people and items. Loose vegetables and other produce that

are billed per pound would presumably rely on sensor tech, but multiple shoppers picking items simultaneously from the same carton would not work with sensors alone. Even pre-packaged or diced vegetables have slight variations in price from one package to another.

Apparel too is particularly hard for computer vision systems to track. Identifying the size (S/M/L) and tracking clothes that are easily folded and tucked away are some of the pain points.

STANDARD COGNITION

Standard Cognition has partnered with Japan's CPG wholesaler Paltac Corporation

AIFI

AiFi reportedly has around 20 retail clients in the pipeline

While startup AiFi promises to utilize existing store infrastructure and a combination of sensors and cameras, Standard Cognition claims to completely do away with sensors, relying solely on machine vision.

Standard Cognition announced a partnership with Paltac Corporation, Japan's largest CPG wholesaler, to outfit 3,000 Japanese stores ahead of the Tokyo Olympics in 2020. AiFi reportedly has around 20 retail clients in the pipeline, including a contract with a major retailer in New York.

In the near term, it comes down to what the cost of deployment and cost of inventory loss due to potential tech glitches would be, and whether a retailer can take on these costs and risks.

BACK OFFICE AUTOMATION

AI is automating administrative work, but the varied nature and formats of data make it a challenging task.

Challenges for automating "back office tasks" can be unique, depending on the industry and the application.

Take clinical trials for instance. Many trials still rely on paper diaries for entering patient data. These diaries are stored digitally, often in difficult-to-search formats, while handwritten clinical notes pose unique challenges for natural language processing algorithms to extract information (accounting for spelling errors, jargon, abbreviations, and missing entries).

Automating auto claims processing, on the other hand, brings a different set of challenges, in this case assessing the damage and drilling down into the root cause.

But different sectors are beginning to adopt ML-based workflow solutions to varying degrees.

Robotic Process Automation (RPA), a loose term for any back office drudge work that is repetitive and can be automated by a bot, has been the subject of much buzz. But, like AI, it's an umbrella term that encompass a wide range of tasks from data entry to compliance to transaction processing to customer onboarding, and more.

Mentions of RPA peaked in earnings calls in Q3'18

Number of mentions of "RPA" or "robotic process automation," by quarter of call

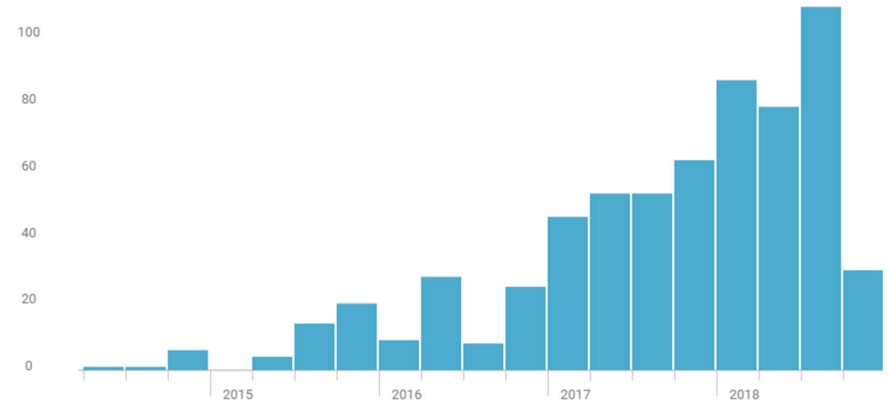

Source: CB Insights earnings transcript search engine

CBINSIGHTS

While not all RPAs are ML-based, many are beginning to integrate image recognition and language processing into their solutions.

WorkFusion, for example, automates back-end operations like Know Your Customer (KYC) and Anti-Money Laundering (AML) processes.

Unicorn UiPath's services have been used by over 700 enterprise clients globally, including DHL, NASA, and HP, across industries ranging from finance to manufacturing to retail.

Automation Anywhere is another unicorn in the RPA space. One of the company's case studies highlights a partnership with a global bank to use machine learning to automate human resource management. An "IQ Bot" extracts information from forms that come in from several countries and in many languages, cleans the data, and then automatically enters it into a human resource management system.

Despite the concept of RPA being around for years, many industries are just beginning to overcome inertia and experiment with newer technologies. In other areas, there's a need for digitization before there can be a layer of predictive analytics.

LANGUAGE TRANSLATION

NLP for language translation is both a challenge and an untapped market opportunity. Big tech companies are pushing the boundaries here.

Machine-based language translation is a huge untapped opportunity with applications in back office automation for multinational corporations, customer support, news & media, and other things.

Baidu recently announced that it's launching new translator earbuds, similar to Google Pixel Buds, which can reportedly translate between 40 different languages in real-time.

Some startups like Unbabel are using human-in-the-loop machine translation systems, with the goal that the feedback loop will train the algorithms to get better over time.

NLP for translation has several challenges. For instance, Chinese natural language processing alone is complex, with 130 spoken dialects and 30 written languages.

A year after Yoshua Bengio, a pioneering researcher in deep learning, published a paper proposing a new architecture for machine translation — a novel way of using neural networks instead of traditional statistical approaches — Google upgraded its own algorithms for the Google Translate Tool.

"This breakthrough will help us provide even more accurate translations for people around the world," CEO Sundar Pichai said in an earnings call in 2016.

Google wanted to move away from its old algorithmic approach of Phrase-Based Machine Translation (PBMT) and proposed a new Google Neural Machine Translation (GNMT) system.

Although different papers had been published on neural machine translation, there were limitations, like the time and computational resources that went into training these models, and failure in translating rare words.

Google suggested improvements to address these issues, and tested its algorithms on English to Chinese, Chinese to English, Spanish to English, among other examples.

Language pair	***Input sentence:***	***Translation (PBMT):***	***Translation (GNMT):***	***Translation (human):***
Chinese->English	2015年到2016年，亚太地区的信用卡交易总额将增加1.7万亿美元，增幅最大，其次是北美地区，增幅将达到1870亿美元。	2015 to 2016 , the total amount of credit card transactions in Asia Pacific will increase by $ 1.7 trillion , the largest increase , followed by North America , growth will reach $ 187 billion .	Total credit card transactions in the Asia-Pacific region will increase by $ 1.7 trillion in 2015-2016, the largest increase, followed by North America with $ 187 billion.	Total credit card transactions in the Asia-Pacific region will increase by $ 1.7 trillion in 2015-2016, the largest increase followed by North America with $ 187 billion.
Chinese->English	100年前，预测引力波的爱因斯坦或许都无法想象人类可以直接观测到引力波。	100 years ago, the prediction of Einstein's gravitational waves probably can not imagine humans can directly observe gravitational waves.	100 years ago, Einstein predicted gravitational waves may not be able to imagine humans can directly observe the gravitational waves.	100 years ago, Einstein who predicted gravitational waves may not be able to imagine that humans can directly observe the gravitational waves.

Source: Google AI Blog

Several research papers have been published on the topic. But the most recent breakthrough comes from Facebook.

According to the paper, “Most research in multilingual NLP focuses on high-resource languages like Chinese, Arabic or major European languages, and is usually limited to a few (most often only two) languages. In contrast, we learn joint sentence representations for 93 different languages, including under-resourced and minority languages.”

ISO3	ISO2	Details			Training corpus size	Tatoeba Error [%]		Tatoeba test set size
		Name	Family	Script		en → xx	xx → en	
hye	hy	Armenian	Armenian	Armenian	6k	59.97	67.79	742
bel	be	Belarusian	Slavic	Cyrillic	5k	31.20	36.50	1000
mya	my	Burmese	Sino-Tibetan	Burmese	2k	n/a	n/a	–
dtp		Central/Kadazan Dusun	Malayo-Polynesian	Latin	1k	92.10	93.50	1000
khm	km	Central Khmer	Khmer	Khmer	625	77.01	81.72	722
cbk		Chavacano	Creole, Romance	Latin	1k	24.20	21.70	1000
kzj		Coastal Kadazan	Malayo-Polynesian	Latin	560	91.60	94.10	1000
cor	kw	Cornish	Celtic	Latin	2k	91.90	93.20	1000
mhr		Eastern Mari	Uralic	Cyrillic	1k	87.70	91.50	1000
ido	io	Ido	constructed	Latin	3k	17.40	15.20	1000
ina	ia	Interlingua	constructed	Latin	9k	5.40	4.10	1000
ile	ie	Interlingue	constructed	Latin	3k	14.70	12.80	1000
gle	ga	Irish	Irish	Latin	732	93.80	95.80	1000
kaz	kk	Kazakh	Turkic	Cyrillic	4k	80.17	82.61	575
lfn		Lingua Franca Nova	constructed	Latin	2k	35.90	35.10	1000
oci	oc	Occitan (post 1500)	Romance	Latin	3k	39.20	38.40	1000
wuu		Wu Chinese	Chinese	Chinese	2k	25.80	25.20	1000
yue		Yue Chinese	Chinese	Chinese	4k	37.00	38.90	1000

Table 2: List of the 18 very low-resource languages included during training of the proposed model, along with their language family, writing system, the resulting similarity error rate on the Tatoeba test set, and the number of sentences in it. Dashes denote language pairs excluded for containing less than 100 test sentences.

Source: https://arxiv.org/pdf/1812.10464.pdf

As big tech companies continue devoting resources to improving translation frameworks, efficiency and language capabilities will improve and adoption will increase across industries.

SYNTHETIC TRAINING DATA

Access to large, labeled datasets is necessary for training AI algorithms. Realistic fake data may solve the bottleneck.

AI algorithms are only as good as the data they are fed, and accessing and labeling this data for different applications is time and capital intensive.

Access to this type of real-world data may not even be feasible.

Consider an autonomous vehicle for instance. Training AVs on dangerous, less frequent situations, such as blinding sun or a pedestrian jumping out from behind parked cars, using real data is hard.

That's where synthetic datasets come in.

In March 2018, Nvidia launched a cloud-based photorealistic simulation for autonomous vehicles called DRIVE Constellation. AVs can drive in virtual reality simulation for billions of miles before hitting the roads — a venture aimed at creating "a safer, more scalable method for bringing self-driving cars to the roads."

Imagine AVs driving through a thunderstorm. Nvidia's solution simulates what data sensors in the car, (like a camera or LiDAR) would generate under these conditions. The synthetic sensor data is fed to a computer which makes decisions as if it were driving on an actual road, sending commands back to the virtual vehicle.

An interesting emerging trend is using AI itself to help generate more "realistic" synthetic images to train AI.

Nvidia, for instance, used generative adversarial networks (GANs) to create fake MRI images with brain tumors.

"Together, these results offer a potential solution to two of the largest challenges facing machine learning in medical imaging, namely the small incidence of pathological findings, and the restrictions around sharing of patient data."

— NVIDIA RESEARCH PAPER

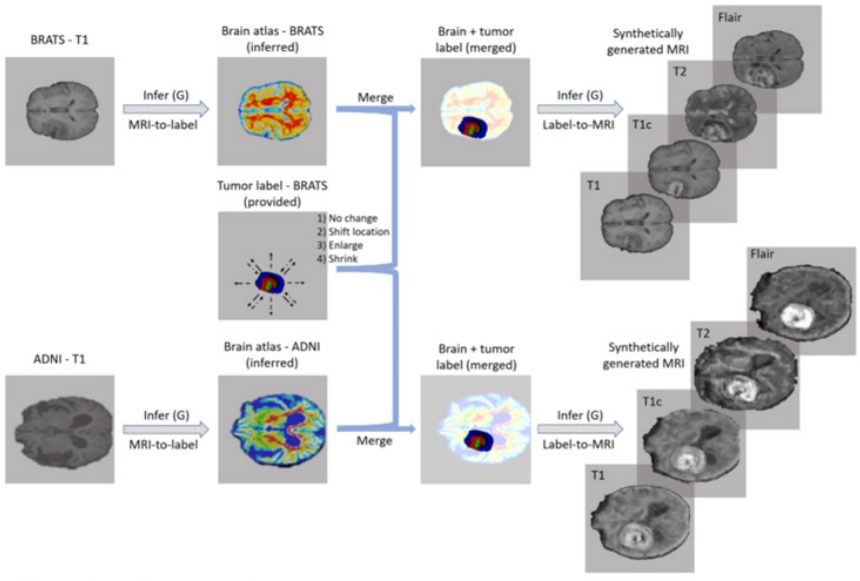

Source: https://arxiv.org/pdf/1807.10225.pdf

GANs are being used to “augment” real world data, meaning AI can be trained with a mix of real world and simulated data to have a larger, more diverse dataset.

Robotics is another field that can greatly benefit from high-fidelity synthetic data.

Consider a simple task of teaching a robot to grasp something. In 2016, Google researchers used 14 robotic arms tasked with learning how to grasp different objects. Data from the failed and successful attempts from all 14 robots were used to train a neural network to help the robots “share their experiences” and predict the outcome of a grasp.

In all, it took 800,000 grasp attempts, “equivalent to about 3000 robot-hours of practice” to “see the beginnings of intelligent reactive behaviors,” according to the research team.

But simulations — having hundreds of virtual robots practice in a virtual environment — can vastly simplify this process.

One of the challenges is creating realistic objects (like making the simulation of an apple or pencil look as close to a real-life objects as possible). In 2017, Google researchers used generative adversarial networks (GANs) to do just that, drastically reducing the amount of real-world data needed to train the robot.

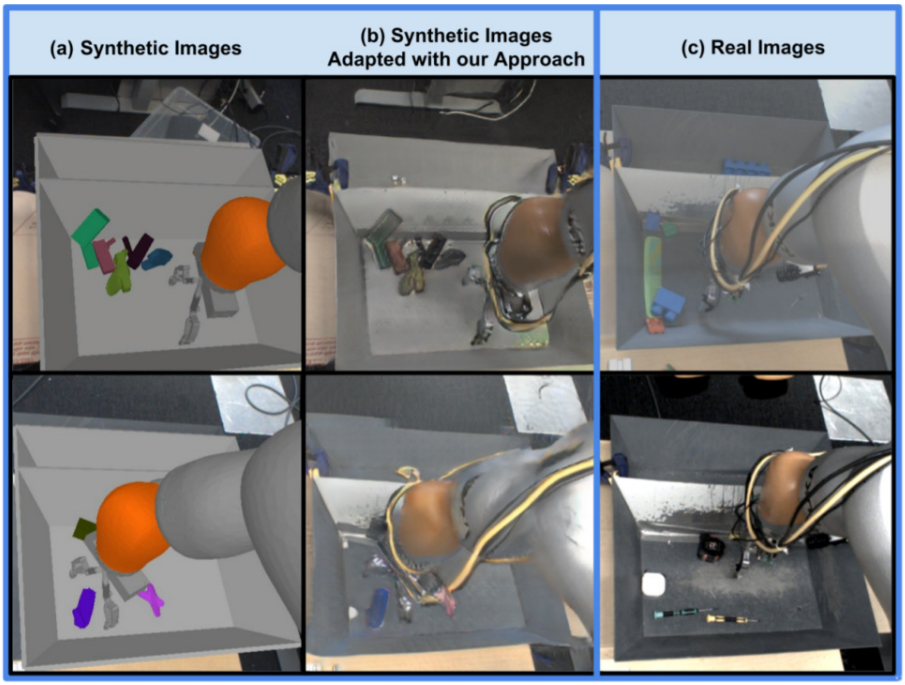

Source: https://arxiv.org/pdf/1709.07857.pdf

Early-stage startups like AI.Reverie are developing simulation platforms to generate datasets for a variety of industries and scenarios.

As the tech scales and synthetic data mimics real-world scenarios more accurately, it will act as a catalyst for smaller companies that don't have access to large datasets.

Threatening

REINFORCEMENT LEARNING

From training algorithms to beat world champions in board games to teaching AI acrobatics, researchers are pushing the boundaries with reinforcement learning. But the need for massive datasets currently limits practical applications.

Reinforcement learning gained media attention when Google DeepMind's AlphaGo defeated a world champion in the complex and strategic Chinese game of Go.

In a nutshell, the point of reinforcement learning is this: What action do you need to take to reach your goal and maximize rewards?

Because of this approach, reinforcement learning has particularly taken off in gaming and robotic simulation.

DeepMind's AlphaGo was initially trained using supervised learning (using data from other human players to train the algorithm) and reinforcement learning (AI playing against itself).

DeepMind later released AlphaGo Zero, which it claimed achieved super-human performance. It was trained purely based on reinforcement learning (playing against itself given just a set of rules).

Recently, researchers at UC Berkeley used computer vision and reinforcement learning to teach algorithms acrobatic skills from YouTube videos. Computer-simulated characters were able to replicate the moves in the videos without the need for manually annotating poses.

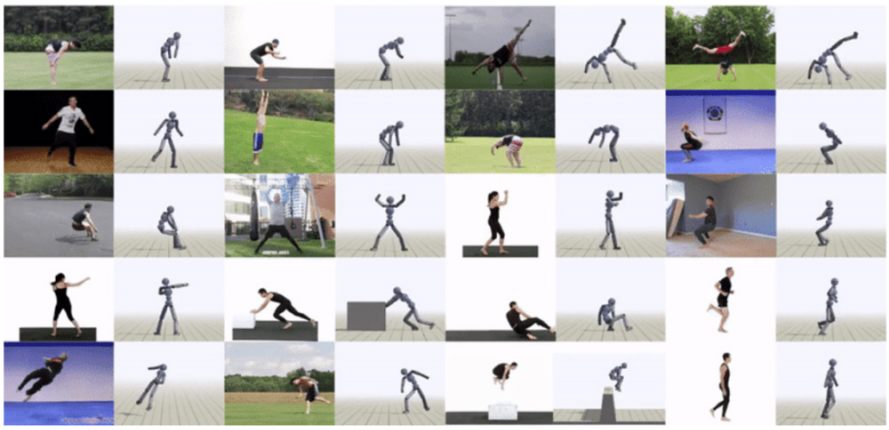

Source: https://bair.berkeley.edu/blog/2018/10/09/sfv/

With reinforcement learning, the simulated characters can apply their skills to new environments. For example, if a man in a YouTube video did a backflip on flat ground, the simulated character can adapt the skill to do a backflip on uneven terrain.

Despite these rapid advances, reinforcement learning adoption hasn't yet taken off because of how much data it requires compared to supervised learning, which is the most prevalent AI paradigm today.

“There’s a rapid fall off as you go down this list [of different approaches to learning] as you think of the economic value created today … Reinforcement Learning is one class of technology where the PR excitement is vastly disproportionate relative to the actual deployments today.”

— ANDREW NG, EMTECH 2017 PRESENTATION

But research into RL applications is increasing. A keyword search in title and abstract of US patent applications shows an uptick in activity in the last 2 years.

US patent applications for reinforcement learning

Number of patents by date of filing, based on keyword searches in title and abstract

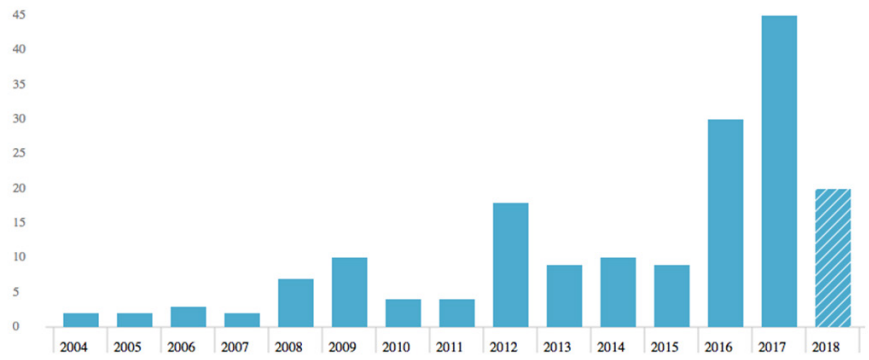

Note: Patterned column(s) may show a decline due to a delay between patent filing and publication

Source: CB Insights patent analytics

CBINSIGHTS

Top applicants include Google, IBM, Alphaics (an AI startup), Mobileye

(acquired by Intel), Microsoft, Adobe, and FANUC.

In earnings calls, Baidu actively discussed reinforcement learning, mentioning it 7 times in its Q1’18 call.

“One highlight in Q1 is that for the first time, we deployed a powerful reinforcement learning based infrastructure that can significantly improve our ability to better match ads to our users and increase clickthrough rates and conversions”

— BAIDU ON A Q1’18 EARNINGS CALL

NETWORK OPTIMIZATION

From facilitating spectrum sharing to monitoring assets and coming up with optimal designs for antenna, AI is beginning to change telecommunication.

Telecommunication network optimization is a set of techniques to improve latency, bandwidth, and design or architecture — anything that augments the flow of data in a favorable way.

For communication service providers, optimization directly translates into better customer experience.

One of the biggest challenges in telecommunications, apart from bandwidth constraints, is network latency. Applications like AR/VR on mobile phones will only optimally function with extremely low lag times.

Apple was granted a patent recently to use machine learning to form "anticipatory networks," which anticipate what action wireless-enabled devices like smartphones may likely perform in the future and download data packets in advance to reduce latency.

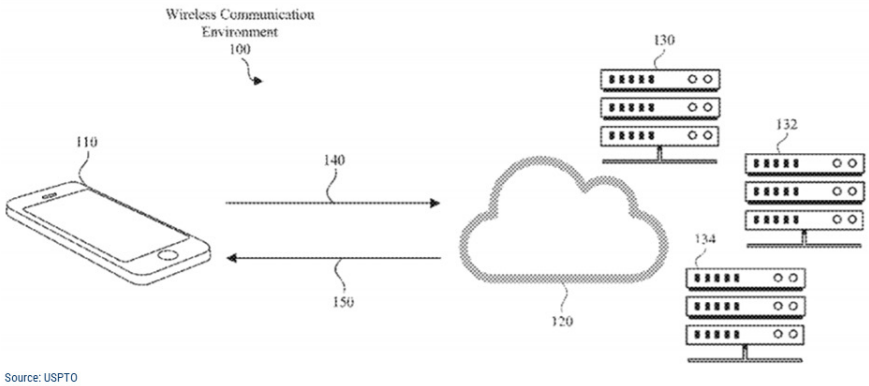

Source: USPTO

Another emerging application of machine learning is in spectrum sharing.

The government licenses certain frequencies of the electromagnetic spectrum to companies like Verizon in an auction.

The Federal Communications Commission (FCC) ruled that the 3.5 to 3.7GHz spectrum will be shared between different users.

This means carriers can dynamically access shared frequencies based on availability. This will allow them to scale bandwidth up and down based on network demand. It will also provide spectrum access to smaller commercial users that don't license a dedicated spectrum of their own.

Parts of the 3.5GHz band is used by the US Navy and other federal agencies. They are given the first tier of access, and if the spectrum is not being used by them, then it goes to tier 2 and tier 3 users.

Companies like Federated Wireless provide Secure Spectrum Access (SAS) to dynamically assign spectrum between different tiers of users and ensure there's no interference with federal signals — and it leverages machine learning to do that.

In 2018, Federated Wireless was granted a patent to use ML to classify radio signals into different categories, such as federal signals, noise signals, and unknown signals. It does this while obscuring features of federal signals (so that hackers never gain access to specific features or weaknesses in military/defense signals).

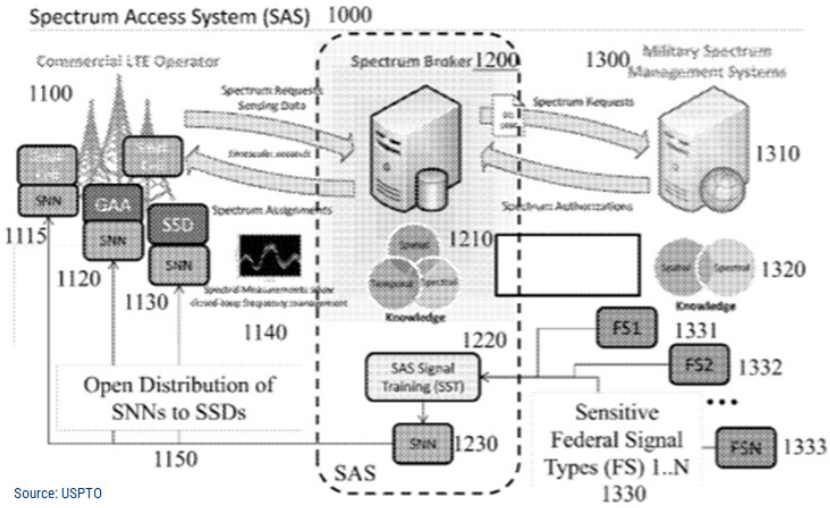

Source: USPTO

DARPA wants to eventually move away from SAS players that facilitate spectrum sharing to an automated ML-based system. To this end, it launched the Spectrum Collaboration Challenge in 2016. Participants in the competition have to use ML to come up with unique ways for radio networks to "autonomously collaborate to dynamically determine how the radio frequency (RF) spectrum should be used moment to moment."

DARPA also launched a Radio Frequency Machine Learning Systems (RFMLS) program in 2017. Similar to the Federated Wireless patent above, DARPA wants to use ML to differentiate between different types of signal, especially spotting malicious signals that intend to hack into end devices (such as IoT devices).

Telecom players are also preparing to integrate AI-based solutions in the next generation of wireless technology, known as 5G.

Samsung acquired AI-based network and service analytics startup Zhilabs in preparation for the 5G era.

Samsung said in a press release that AI software will be used to “analyze user traffic, classify applications being used, and improve overall service quality.”

Qualcomm sees AI edge computing as a crucial component of its 5G plans (edge computing reduces bandwidth constraints and frequent communications with the cloud — a main focus area for 5G).

Early research papers are also emerging exploring the use of neural nets to come up with the most optimal design for antenna in telecommunication networks.

AUTONOMOUS VEHICLES

Despite a substantial market opportunity for autonomous vehicles, the timeline for full autonomy is still unclear.

A number of big tech companies and startups are competing intensely in the autonomous vehicles space.

Google has made a name for itself in the auto space. Its self-driving project Waymo is the first autonomous vehicle developer to deploy a commercial fleet of AVs.

Investors remain confident in companies developing the full autonomous driving stack, pouring hundreds of millions of dollars into GM's Cruise Automation ($750M from Honda in October 2018 and $900M from SoftBank in May prior) and Zoox ($500M in July 2018). Other startups here include Drive.ai, Pony.ai, and Nuro.

China, in particular, has ramped up its AV efforts. The Chinese science ministry announced last year that the nation's first wave of open AI platforms will rely heavily on Baidu for autonomous driving.

In April 2017, Baidu announced a one-of-a-kind open platform — Apollo — for autonomous driving solutions, roping in partners from across the globe.

As with other open-source platforms, the idea is to accelerate AI and autonomous driving research by opening it up to contributions from other players in the ecosystem. Making the source code available to everyone allows companies to build off of existing research instead of starting from scratch.

Baidu builds global autonomous driving ecosystem

Select international partners of Baidu's Apollo project

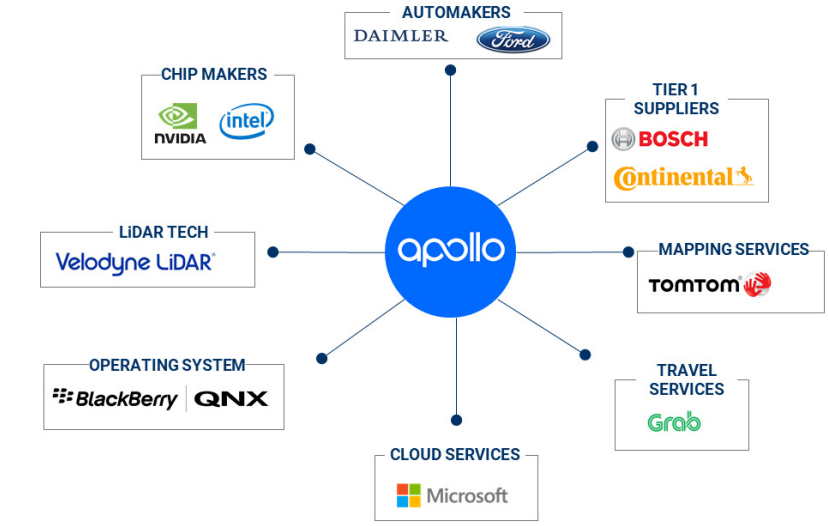

Partnerships source: Apollo.auto

CBINSIGHTS

Alibaba also recently conducted test drives of its autonomous vehicle. But interestingly, just over a year ago, Alibaba was skeptical about the long-term commercial opportunity of autonomous vehicles, mentioning in an earnings call that "nobody has figured out the long-term economic model for this, but people are doing it because there is some very interesting artificial intelligence-related technology" involved in building autonomous vehicles.

Even with hesitation surrounding the future of the technology, automakers are still working full steam ahead. The market is projected to reach roughly $80B by 2025.

Some applications could see earlier adoption of fully self-driving vehicles, such as logistics and fulfillment.

Autonomous last-mile delivery is making news

Number of news mentions for “autonomous” and “last mile”

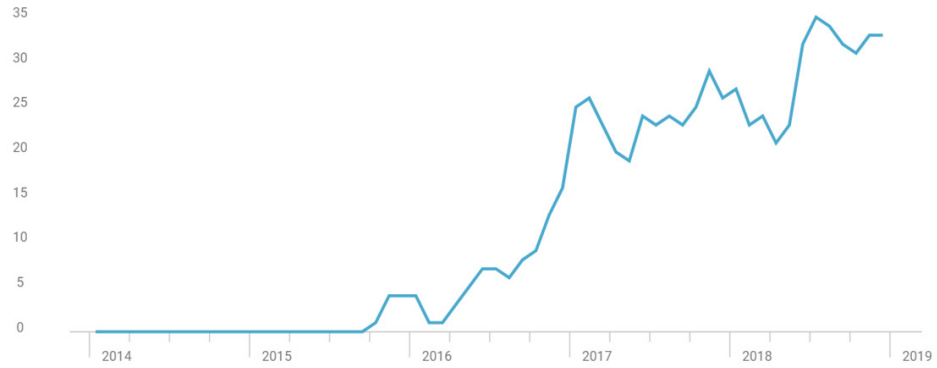

Source: CB Insights News Trends tool

CBINSIGHTS

Autonomous logistics — specifically autonomous last-mile delivery — is top-of-mind for retailers and fulfillment companies, and may be the first area where we see full autonomy. Self-driving vehicles could help tackle the costly and arduous challenge of delivering goods at the last mile, which can add up to nearly a third of an item’s total delivery cost.

States like Arizona which have liberal laws for autonomous vehicle deployment are emerging as test beds. In June 2018, robotics startup Nuro partnered with Kroger, one of the largest brick-and-mortar grocers in the US, to deliver groceries. Nuro is designed to drive on neighborhood roads, not just sidewalks like other delivery robot and vehicle prototypes that have been developed.

In the restaurant space, pizza companies like Domino’s and Pizza Hut have been at the forefront of testing out autonomous vehicles. Ford is piloting autonomous delivery in Miami with pizza, groceries, and other goods. The OEM partnered with over 70 businesses, including Domino’s, in early 2018.

CROP MONITORING

Three types of crop monitoring are taking off in agriculture: On-ground, aerial, and geospatial.

The precision agriculture drone market is expected to reach $2.9B in 2021.

Drones can map the field for farmers, monitor moisture content using thermal imaging, and identify pest infested crops and spray pesticides.

Startups are focusing on adding a layer of analytics to data captured by 3rd party drones.

Taranis, for example, uses 3rd party Cessna airplanes to do this. Taranis also acquired agtech-AI startup Mavrx Imaging last year, which was developing ultra high resolution imaging tech to scout and monitor fields.

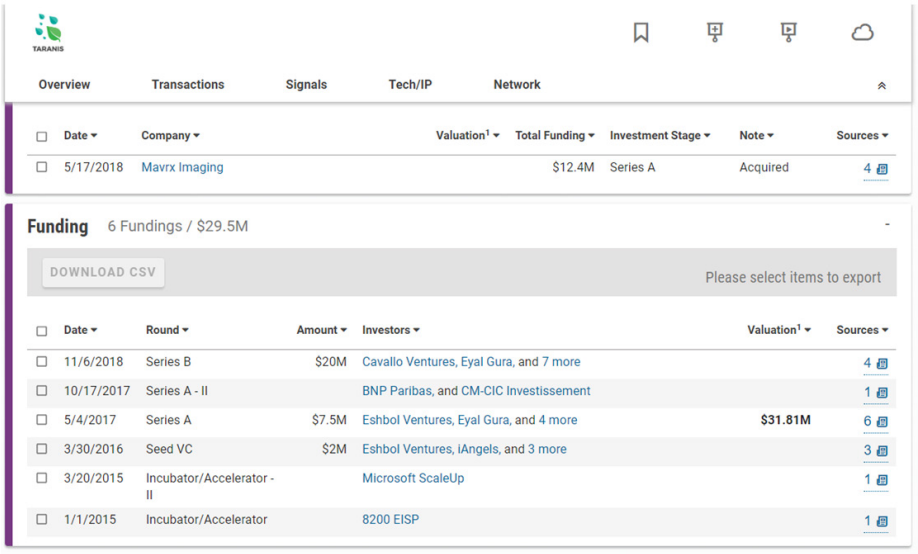

Source: CB Insights

Taranis uses AI to stitch together images of the field and also to identify potential issues with crops. John Deere, a farming equipment manufacturer, tapped the startup along with a few others, to collaborate on potential solutions for John Deere.

Deere has been reinventing itself with AI. It bought Blue River Technology — an agricultural equipment company leveraging computer vision — for $300M+. Among other things, Blue River was working on "smart weeding" and "see-and-spray" solutions.

This type of individual crop monitoring can become a major disruptor for the agricultural pesticide industry. If on-the-ground farming equipment gets smarter with computer vision and sprays only individual crops as needed, it will reduce the demand for non-selective weed killers that kill everything in the vicinity. Precision spraying would also mean a reduction in the amount of herbicide and pesticide used.

Beyond the field, using computer vision to analyze satellite images provides a macro-level understanding of agricultural practices.

Geo-spatial data can provide information on crop distribution patterns across the globe and the impact of weather changes on agriculture.

Cargill invested in Descartes Labs, which uses satellite data to develop a forecasting model for crops like soybean and corn. This application of computer vision has also piqued the interest of commodities traders and government agencies. DARPA is working with Descartes to forecast food security.

Transitory

CYBER THREAT HUNTING

Reacting to cyber attacks is no longer enough. Proactively "hunting" for threats using machine learning is gaining momentum in cybersecurity.

Advancements in computing power and algorithms are turning previously theoretical hacks into real security problems.

According to the Breach Level Index, a global database of public data breaches, 4.5B data records were compromised worldwide in H1'18 (for reference, the figure was 2.6B for all of 2017).

Rank	Organization Breached	Records Breached	Date of Breach	Type of Breach	Source of Breach	Location	Industry	Risk Score
1	Facebook	2,200,000,000	04/04/18	Identity Theft	Malicious Outsider	United States	Social Media	10.0
2	Exactis	340,000,000	06/01/18	Identity Theft	Accidental Loss	United States	Other	9.1
3	Under Armour	150,000,000	02/01/18	Account Access	Malicious Outsider	United States	Retail	9.1
4	Twitter	336,000,000	05/03/18	Financial Access	Accidental Loss	United States	Social Media	9.0
5	Firebase (Google)	100,000,000	06/20/18	Identity Theft	Accidental Loss	United States	Technology	8.6

Source: Breach Level Index top H1'18 breaches

Unlike other industrial applications of AI, cyber-defense is a cat-and-mouse game between hackers and security personnel, both leveraging advances in machine learning to up their game and keep ahead of the other.

Threat hunting, as the name suggests, is the practice of proactively seeking out malicious activity instead of merely reacting to alerts or a breach after it has occured.

Hunting begins with a hypothesis on potential weaknesses in the network, and manual and automated tools to test out the hypothesis in a continuous, iterative process. The sheer volume of data in cybersecurity makes machine learning an inseparable part of the process.

A quick search on Linkedin for "threat hunters" shows 70+ job listings in the United States from organizations such as Microsoft, Raytheon, Verizon, Booz Allen Hamilton, and Dow Jones.

While this reflects an emerging demand for threat hunters across diverse business types, it also indicates that the title itself is still niche.

"Results from the SANS 2018 Threat Hunting Survey show that, for many organizations, hunting is still new and poorly defined from a process and organizational standpoint … The survey of 600 respondents reveals that most organizations that are hunting tend to be larger enterprises or those that have been heavily targeted in the past."

- SANS 2018 SURVEY SPONSORED BY IBM

As the SANS 2018 survey suggests, the stakes are higher for larger enterprises whose differentiating factor is their access to a treasure trove of data.

Amazon, for instance, faces mounting pressure from AWS customers to secure the cloud. Wrongfully configured AWS servers have resulted in data breaches at customers like Verizon, WWE, Dow Jones, and Accenture.

Amazon acquired threat hunting startup Sqrrl to develop a new product for hunting hackers on AWS clients' accounts.

Cylance, another AI startup with a focus on threat hunting, was acquired by Blackberry last year.

The more spread out a network becomes the more vulnerable it becomes.

Threat hunting is likely to gain further traction, however it does come with its own set of challenges, such as dealing with an ever-changing, dynamic environment and reducing false positives.

CONVERSATIONAL AI

For many enterprises, chatbots became synonymous with AI — but the promise isn't keeping up with the reality.

Recently, Google was in hot water over its conversational AI feature, Duplex.

Duplex can make phone calls and reservations on behalf of the user, but communicates like a real human (complete with "umms" and pauses). It sparked ethical concerns over whether or not Duplex needs to identify itself as a conversational agent when speaking to real people.

Google added Duplex to its new phone, Pixel 3. It has turned the Pixel 3 into an AI powerhouse, including a "screen call" option that allows the Google Assistant to screen for spam callers.

Google has been applying to patent the interactions between two conversational agents since 2014. The most recent application, "Conversational Agent Response Determined Using A Sentiment," was filed in April 2018.

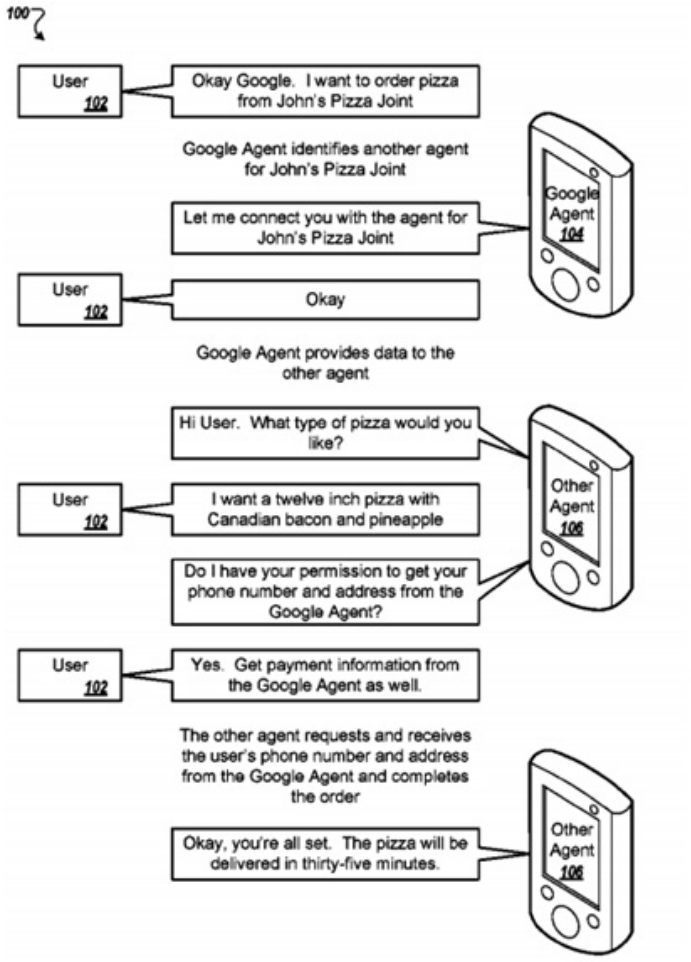

Source: USPTO

Despite FAMGA and China's big tech companies (Baidu, Alibaba, and Tencent) focusing heavily on this space, conversational agents — both voice- and text-based — are more feasible in some applications than others.

One of the most widespread applications of chatbots is in customer service. Bots form the first layer of interaction with the user (note: not all bots use natural language processing) and hand off queries to a human based on the level of complexity.

This is still challenging for applications like health and insurance, where triaging (gauging the urgency of a situation) is complex.

Similarly, shopping through voice-based conversations alone, without a visual cue, is challenging.

Although analysts and CPG brands, from Sephora and Nestle to Capgemini, have talked up voice shopping as the next big thing in retail, it hasn't taken off. With the exception of reordering specific items, it fails to provide key customer experiences that drive online commerce.

Mental healthcare is another area where chatbots seem like a potentially disruptive force.

High costs of mental health therapy and the appeal of round-the-clock availability is giving rise to a new era of AI-based mental health bots. Early-stage startups are focused on using cognitive behavioral therapy — changing negative thoughts and behaviors — as a conversational extension of the many mood tracking and digital diary wellness apps in the market.

But mental health is a spectrum. There is variability in symptoms, subjectivity in analysis, and it requires a high level of emotional cognition and human-to-human interaction.

This makes areas like mental healthcare — despite the upside of cost and accessibility — a particularly hard task for algorithms.

DRUG DISCOVERY

With AI biotech startups emerging, traditional pharma companies are looking to AI SaaS startups for innovative solutions to the long drug discovery cycle.

In May 2018, Pfizer entered into a strategic partnership with XtalPi — an AI startup backed by tech giants like Tencent and Google — to predict pharmaceutical properties of small molecules and develop "computation-based rational drug design."

But Pfizer is not alone.

Top pharmaceutical companies like Novartis, Sanofi, GlaxoSmithKline, Amgen, and Merck have all announced partnerships in recent months with AI startups to discover new drug candidates for a range of diseases from oncology and cardiology.

"The biggest opportunity where we are still in the early stage is to use deep learning and artificial intelligence to identify completely new indications, completely new medicines."

— BRUNO STRIGINI, FORMER CEO OF NOVARTIS ONCOLOGY

Interest in the space is driving the number of equity deals to AI drug

discovery startups: 20 as of Q2’18, equal to all of 2017.

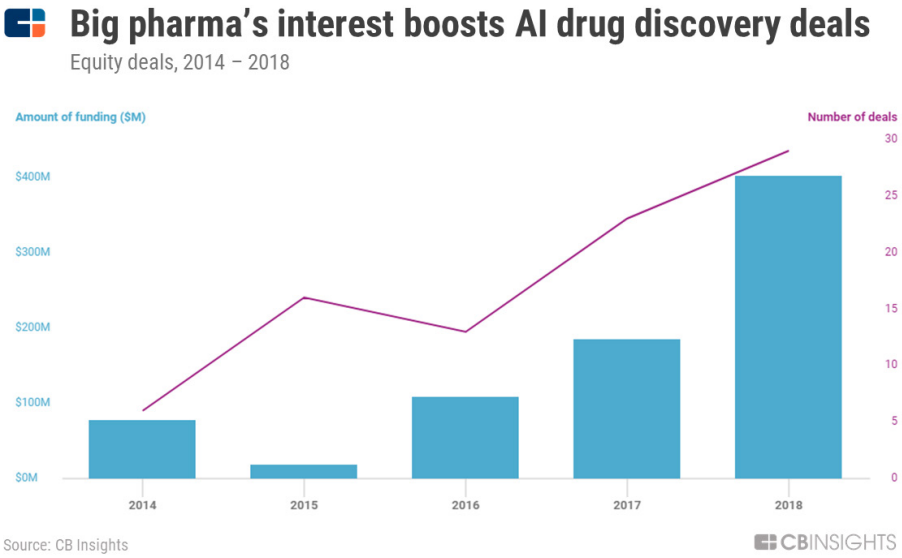

While biotech AI companies like Recursion Pharmaceuticals are investing in both AI and drug R&D, traditional pharma companies are partnering with AI SaaS startups.

Although many of these startups are still in the early stages of funding,

they already boast a roster of pharma clients.

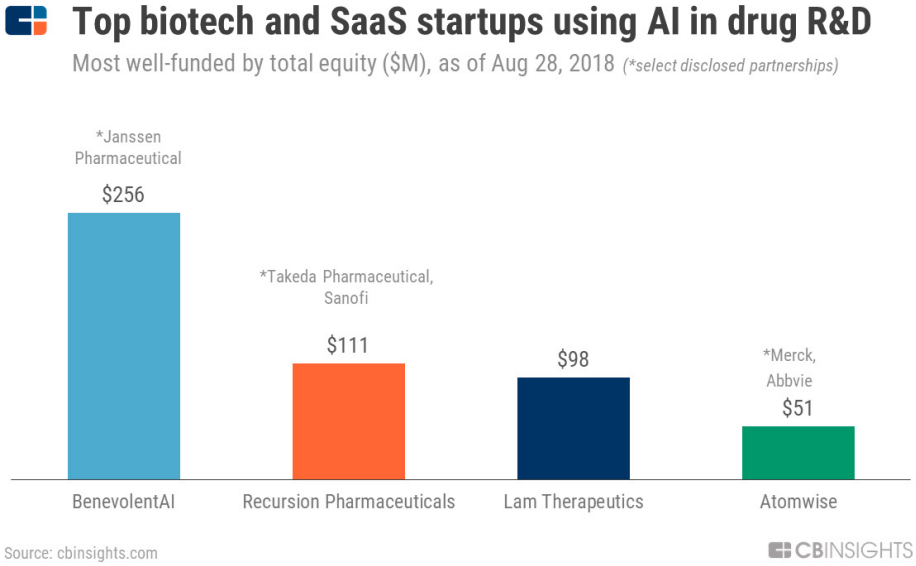

There are few measurable metrics of success in the drug formulation phase, but pharma companies are betting millions of dollars on AI algorithms to discover novel therapeutic candidates and transform the drawn-out drug discovery process.

WHERE IS ALL THIS DATA FROM?

The CB Insights platform has the underlying data included in this report

CLICK HERE TO SIGN UP FOR FREE